

PREPARED FOR SUBMISSION TO JCAP

Optimal bispectrum constraints on single-field models of inflation

Gemma J. Anderson, Donough Regan and David Seery

Astronomy Centre, University of Sussex,
Falmer, Brighton, BN1 9QH, UK

E-mail: G.Anderson@sussex.ac.uk, D.Regan@sussex.ac.uk, D.Seery@sussex.ac.uk

Abstract. We use WMAP 9-year bispectrum data to constrain the free parameters of an ‘effective field theory’ describing fluctuations in single-field inflation. The Lagrangian of the theory contains a finite number of operators associated with unknown mass scales. Each operator produces a fixed bispectrum shape, which we decompose into partial waves in order to construct a likelihood function. Based on this likelihood we are able to constrain four linearly independent combinations of the mass scales. As an example of our framework we specialize our results to the case of ‘Dirac–Born–Infeld’ and ‘ghost’ inflation and obtain the posterior probability for each model, which in Bayesian schemes is a useful tool for model comparison. Our results suggest that DBI-like models with two or more free parameters are disfavoured by the data by comparison with single-parameter models in the same class.

Contents

1	Introduction	1
2	Overview of the effective field theory of inflation	3
3	Calculation of the bispectrum	6
4	Estimating the EFT mass scales	8
4.1	How many independent shapes?	12
4.2	Results	15
5	Constraints on models	16
6	Model comparison using the bispectrum	20
7	Discussion and conclusions	24
A	Three-point functions for the EFT operators	25

1 Introduction

Successive microwave-background surveys have accumulated some evidence for the inflationary paradigm, in which structure in the universe was seeded by quantum fluctuations during an epoch preceding the hot, dense phase where nucleosynthesis occurred [1, 2]. But despite broad support for the overall framework, attempts to identify the precise degrees of freedom whose quantum fluctuations were relevant have met with less success. Whatever microphysics underlay the putative inflationary epoch remains mysterious.

In scattering experiments, an abundance of observables—including, among others, branching ratios, decay rates, and differential dependence on energy or angles—allow indirect access to microphysical information through reconstruction of the correlation functions, or ‘ n -point functions’. These measure interference between quantum fluctuations and encode information about the dynamics of the theory. It is the rich information which can be obtained from reconstruction of the correlation functions which makes measurements in particle physics so constraining.

In cosmology our observables are more limited and so is the degree to which the n -point functions can be reconstructed. Over a narrow range of scales, the n -point functions of the cosmic microwave background (‘CMB’) anisotropies are sensitive to the n -point functions of the primordial ‘curvature perturbation’, which is a calculable, model-dependent mix of the fluctuations imprinted on the light fields of the inflationary epoch. This correspondence has been used for many years to place restrictions on the inflationary model space from measurements of the CMB temperature and polarization two-point functions. But if a *three*-point function of the CMB anisotropies could be measured it would provide access to more nuanced and discriminating microphysical information. Ideally we would like to observe systematic relationships between the n -point functions which would point clearly to a quantum mechanical origin for the fluctuations. This is important because it is unclear whether we could ever

rule out a non-quantum origin (perhaps associated with new but non-inflationary physics at early times) using only the two-point function.

Measurements of the CMB temperature anisotropy have now reached sufficient accuracy that it is feasible to estimate the three-point temperature autocorrelation function. The most precise constraints come from the Planck2013 dataset [1]. But despite the quality of the measurements, the signal-to-noise for any particular combination of wavenumbers is still too low to allow the three-point function to be reconstructed directly. Instead, measurements are made by picking an Ansatz or ‘template’ for the way in which the correlations change with wavenumber. By comparing this template with the CMB data over many different combinations of wavenumber it is possible to attain reasonable signal-to-noise. This comparison carries a considerable computational burden, so constraints from the data are typically reported as amplitudes for just a handful of well-known templates, such as the ‘local’, ‘equilateral’ and ‘orthogonal’ shapes. These amplitudes are often written $\hat{f}_{\text{NL}}^{\text{local}}$, $\hat{f}_{\text{NL}}^{\text{equi}}$, $\hat{f}_{\text{NL}}^{\text{ortho}}$, and so on.¹

A specific inflationary model will be characterized by a number N_λ of adjustable parameters λ_i , $1 \leq i \leq N_\lambda$. These may include Lagrangian parameters which are analogues of masses and couplings, but in multiple-field models may also include a specification of the initial conditions in field-space. To apply constraints from $\hat{f}_{\text{NL}}^{\text{local}}$, $\hat{f}_{\text{NL}}^{\text{equi}}$, $\hat{f}_{\text{NL}}^{\text{ortho}}$, \dots , to such a model its three-point function must be computed and projected on to each of these templates. This generates predictions for each of the amplitudes $f_{\text{NL}}^{\text{local}}(\lambda_i)$, $f_{\text{NL}}^{\text{equi}}(\lambda_i)$, $f_{\text{NL}}^{\text{ortho}}(\lambda_i)$, \dots . The results obtained by a microwave background survey can then be converted into constraints on the underlying parameters λ_i .

This approach is perfectly reasonable, but there are reasons to expect that it may not be optimal. First, if the set of templates does not cover the entire range of three-point correlations which can be produced by adjusting the parameters λ_i then we are not making efficient use of the data: we should measure the amplitude of more templates in order to obtain better constraints. But, as many authors have pointed out, it is not clear *a priori* how large a range of templates is required, or how they should be chosen.

Second, if our templates are chosen injudiciously then there will come a point of diminishing returns at which no new information is gained because the shapes we are fitting are strongly correlated with shapes which have been tried before. This is a reflection of a more general problem: the error bars reported for any set of amplitudes will typically be correlated, with the correlation described by some covariance matrix. Without knowledge of these covariances we risk underestimating the uncertainties associated with our reconstruction of the parameters λ_i .

In this paper we take a different approach. We investigate the construction of maximum-likelihood estimators for the Lagrangian parameters λ_i directly from the data. (Because noise maps for the Planck2013 data release are not yet available, we use the WMAP 9-year dataset.) To decide which templates to use, we catalogue the different types of correlation which can be produced in a well-specified class of models: those whose fluctuations are described by the effective field theory of inflation [3]. We construct the Fisher matrix associated with these correlations and use it to determine the principal directions whose amplitudes can be measured efficiently. We account for the covariance between measurements of these amplitudes and use them to place constraints on the underlying Lagrangian parameters.

Summary.—In §2 we briefly review the effective field theory approach to single-field inflation and catalogue the operators arising from a general single-field action. In §3 we discuss the

¹Here and throughout the remainder of the paper we distinguish quantities estimated from data by a hat.

calculation of bispectra corresponding to these operators, and point out a number of subtleties which must be borne in mind when interpreting our results. In §4 we assemble the formalism which is used to extract constraints from the CMB map: in §4.1 we construct the Fisher matrix and use it to determine the principal directions which can be constrained efficiently, and in §4.2 we report our measurements of their amplitudes from the 9-year WMAP dataset. §5 translates these general constraints into the language of specific models, and §6 uses the framework of Bayesian model comparison to gain some qualitative information regarding the type of model favoured by the data. We conclude in §7. A short appendix tabulates the three-point functions used in the main text.

Notation.—We use units in which $c = \hbar = 1$, and define the reduced Planck mass M_P to be $M_P^{-2} = 8\pi G$. Our index and summation conventions are explained in the main text.

2 Overview of the effective field theory of inflation

In this paper we focus on single-field models of inflation which terminate in a unique minimum, which we refer to as the ‘reheating minimum’. In multiple-field models there are complications associated with our freedom to set initial conditions. These determine the average field-space trajectory followed by the region of the universe we choose to study. In a single-field model there is a unique trajectory which terminates in the reheating minimum.

In both single- and multiple-field cases it is quantum fluctuations around this average field-space trajectory which are inherited by the large-scale density perturbation, but where there is no unique trajectory the calculation of these fluctuations is a serious computational challenge. Their evolution must be followed until an ‘adiabatic limit’ has been reached, at which all isocurvature modes become exhausted [4–8]. Normally this will require numerical methods. In contrast, the fluctuations produced in single-field inflation—or, more precisely, ‘single-clock’ inflation—typically do not evolve and can be computed analytically under certain circumstances. Below, we discuss the precise conditions which are required.

Model parametrization.—Our aim is to estimate the Lagrangian parameters which characterize a single-field inflationary model. How many such parameters are needed? The answer depends on the range of behaviour which we allow. Cheung et al. gave an argument based on nonlinearly realized Lorentz invariance which, under certain conditions, constrains the possible three-body interactions between scalar perturbations on a smooth inflationary background [3]. This is the ‘effective field theory of inflation’. In this section we briefly review their construction.

The effective field theory is not used to describe the background cosmology, but only fluctuations around it. Therefore it is agnostic regarding the precise mechanism of inflation. The background is assumed to be described by a Robertson–Walker metric

$$ds^2 = -dt^2 + a^2(t) d\mathbf{x}^2, \quad (2.1)$$

where $a(t)$ is the scale factor, t is cosmic time and $H(t) = \dot{a}/a$ is the Hubble rate. Since the background is evolving it spontaneously breaks time-translation invariance (and therefore manifest Lorentz invariance), but because the spatial slices are homogeneous and isotropic the background remains manifestly invariant under spatial coordinate transformations. We will use the terminology ‘coordinate transformations’ and ‘diffeomorphisms’ interchangeably.

Knowledge of the background evolution is equivalent to specifying $H(t)$ as a smooth function of t . The condition that the universe is ‘single-clock’ is that a coordinate system

exists in which only the metric carries fluctuations; in this coordinate system all fields needed to describe the matter sector are homogeneous, depending only on the time t . By analogy with similar constructions in particle physics, Cheung et al. called this coordinate system the *unitary gauge*. Where the matter sector is described by a single scalar field ϕ it corresponds to the gauge where fluctuations $\delta\phi$ vanish, but this is not necessary.

To describe dynamics we require a Lagrangian. A Lagrangian which is manifestly invariant under the unbroken (linearly-realized) group of purely spatial coordinate transformations will be a function F which transforms as a scalar under these diffeomorphisms. Cheung et al. argued that the most general such Lagrangian could be constructed as a scalar function of the metric and the intrinsic and extrinsic curvature tensors on the spatial slices, together with their covariant derivatives [3]. These may appear in arbitrary combinations with t and the metric function g^{00} , which are both invariant under spatial coordinate transformations. Therefore,

$$S_{\text{gen}} = \int d^4x \sqrt{-g} F(R_{\mu\nu\rho\sigma}, K_{\mu\nu}, \nabla_\mu, g^{00}, t). \quad (2.2)$$

By itself, this Lagrangian can describe fluctuations around any cosmological background with linearly-realized spatial diffeomorphism invariance. Specializing it to the background $H(t)$ fixes the background and linear terms,

$$S = \int d^4x \sqrt{-g} \left(\frac{M_{\text{P}}^2}{2} R + M_{\text{P}}^2 \dot{H} g^{00} - M_{\text{P}}^2 (3H^2 + \dot{H}) + \sum_{n \geq 2} F_n(\delta R_{\mu\nu\rho\sigma}, \delta K_{\mu\nu}, \nabla_\mu, \delta g^{00}, t) \right), \quad (2.3)$$

where $\delta R_{\mu\nu\rho\sigma}$ and $\delta K_{\mu\nu}$ are, respectively, perturbations in the intrinsic and extrinsic curvature tensors, and $\delta g^{00} = g^{00} + 1$ is the perturbation in the time-time metric function or ‘lapse’. The arbitrary functions F_n are homogeneous polynomials of order n , and therefore the leading correction to the first three terms appearing in (2.3) is quadratic.

We have not yet made use of the requirement that the full theory is invariant under time reparametrizations, $t \rightarrow t' = t + \xi(\mathbf{x})$, where the translation ξ may be a function of position.² On an expanding cosmological background this symmetry is spontaneously broken. Nevertheless, once a choice of spatially-invariant operators has been made in Eq. (2.3), the broken time-translation symmetry is strong enough to fix the interactions of one scalar mode. To determine these interactions we construct a new action by formally performing a time translation $t \rightarrow t' = t - \pi$. If we promote π to a dynamical field which shifts linearly under time translations (that is, $\pi \rightarrow \pi' = \pi - \xi$ when $t \rightarrow t' = t + \xi$) then the total action becomes manifestly invariant. The field π represents a scalar degree of freedom in the system, but its interactions are fixed uniquely by the combination of tensors appearing in the F_n , the background cosmology $H(t)$, and the time translation symmetry [3, 9–12]

For this formalism to be useful it must be possible to calculate each amplitude of interest using states which contain no more than a handful of π particles, or π -lines in diagrammatic terms. This is not generally true. But if all background fields are time-independent then rigid time translations $t \mapsto t' = t + \xi$ (with ξ a constant) are a *global* symmetry of the theory, no matter what transformation law we ascribe to π . Therefore π must behave as a Goldstone boson: where it appears in the action it must be accompanied by at least one derivative. In a process which takes place at a well-defined characteristic energy scale E , each derivative will translate to a power of E . The justification for neglecting diagrams which contain a large

²An arbitrary action of the form (2.3) can describe theories with this symmetry, in addition to others which do not.

number of π -lines is then the same as any effective field theory of Goldstone modes, enabling a perturbative expansion in powers of E/M where M is some large mass scale characterizing the strength of the interactions.

For applications to inflation the background fields are not constant but slowly varying, so rigid time translations are only an approximate symmetry. Therefore terms involving undifferentiated powers of π may appear in the action, although suppressed by dimensionless factors which measure the degree to which the global symmetry is broken. These generate effects which are unaccompanied by powers of E/M and therefore may be important at all energies.³ However, provided the approximate symmetry is sufficiently good that corrections to it are at least as small as the first neglected power of E/M it is still possible to carry out a consistent calculation. During inflation we are interested in the type of correlations induced by each operator between modes of the quantized field near the epoch of Hubble exit, so the scale E will be of order the Hubble scale H .

At sufficiently high energies $E > E_{\text{mix}}$ the Goldstone mode decouples from the remaining degrees of freedom in $\delta R_{\mu\nu\rho\sigma}$ and $\delta K_{\mu\nu}$. (The notation ‘ E_{mix} ’ was introduced by Cheung et al. [3], who emphasized that below E_{mix} the mixing with gravitational degrees of freedom cannot be ignored.) If the decoupling scale E_{mix} is at least modestly smaller than $E = H$ then it is possible to study how each operator generates correlations without including gravitational fluctuations. In this paper we will work exclusively in the decoupling limit. With this assumption, Bartolo et al. [13, 14] gave an effective action up to cubic terms,

$$\begin{aligned}
S_{\text{EFT}} = \int d^4x \sqrt{-g} \Big\{ & M_{\text{P}}^2 \dot{H} (\partial_\mu \pi)^2 + 2M_2^4 \left[\dot{\pi}^2 - \dot{\pi} \frac{(\partial\pi)^2}{a^2} \right] - \frac{4}{3} M_3^4 \dot{\pi}^3 \\
& - \frac{\bar{M}_1^3}{2a^2} \left[-2H(\partial\pi)^2 + \frac{(\partial\pi)^2 \partial^2 \pi}{a^2} \right] \\
& - \frac{\bar{M}_2^2}{2a^4} \left[(\partial^2 \pi)(\partial^2 \pi) + H(\partial^2 \pi)(\partial\pi)^2 + 2\dot{\pi} \partial^2 \partial_j \pi \partial_j \pi \right] \\
& - \frac{\bar{M}_3^2}{2a^4} \left[(\partial^2 \pi)(\partial^2 \pi) + 2H \partial^2 \pi (\partial\pi)^2 + 2\dot{\pi} \partial^2 \partial_j \pi \partial_j \pi \right] \\
& - \frac{2\bar{M}_4^3}{3a^2} \dot{\pi}^2 \partial^2 \pi + \frac{\bar{M}_5^2}{3a^4} \dot{\pi} (\partial^2 \pi)^2 + \frac{\bar{M}_6^2}{3a^4} \dot{\pi} (\partial_i \partial_j \pi)^2 - \frac{\bar{M}_7}{3! \cdot a^6} (\partial^2 \pi)^3 \\
& - \frac{\bar{M}_8}{3! \cdot a^6} \partial^2 \pi (\partial_i \partial_j \pi)^2 - \frac{\bar{M}_9}{3! \cdot a^6} \partial_i \partial_j \pi \partial_j \partial_k \pi \partial_k \partial_i \pi \Big\}.
\end{aligned} \tag{2.4}$$

Our notation has been chosen to match Refs. [13, 14]. The mass scales M_2 , M_3 and $\bar{M}_1, \dots, \bar{M}_9$ characterize the model under consideration.⁴ Terms decorated with a bar are associated with operators involving the extrinsic curvature $\delta K_{\mu\nu}$, whereas unbarred terms correspond to powers of δg^{00} . In writing Eq. (2.4), Bartolo et al. did not include all possible operators: they neglected higher-derivative operators containing derivatives of the form $\nabla_\mu \delta g^{00}$ and $\nabla_\lambda K_{\mu\nu}$, and from the lowest-derivative combinations for each M_i or \bar{M}_i they retained only terms which gave a parametrically large contribution to the three-point function. We can expect

³It is these terms which cause superhorizon evolution of the perturbations in multiple-field models. Their importance at all scales is reflected in the fact that they remain relevant even when k/aH is very soft.

⁴To aid intuition, the powers of the M_i and \bar{M}_i appearing in Eq. (2.4) have been chosen so that the M_i and \bar{M}_i all have dimensions of mass when using natural units in which $c = \hbar = 1$. In some cases this means that positive integer powers of masses appear, such as M_3^4 , which can only be positive if M_3 is real. In reality there is an undetermined sign which we are suppressing, so that M_3^4 should be regarded as an object which can be either positive or negative. The associated mass scale is $|M_3^4|^{1/4}$.

the higher-derivative operators to be small provided the mass scales M_i , \bar{M}_i are sufficiently large, which is already the condition that the EFT is predictive. Therefore, although (2.4) does not represent the most general set of interactions, it is reasonable to speculate that it may approximate the most general set of *observable* interactions for a smooth background $H(t)$. In this paper we only consider backgrounds which satisfy this smoothness requirement. The properties of fluctuations over backgrounds which are not sufficiently smooth require a separate analysis; for example, see Refs. [15, 16].

When is the decoupling approximation valid? Estimates for the scale E_{mix} were given by Cheung et al. [3], but strictly this scale can be determined only when the M_i and \bar{M}_i are known and therefore it must be checked *a posteriori*. As an example, in canonical single-field inflation, Cheung et al. argued that $E_{\text{mix}} \sim \epsilon^{1/2} H$, where $\epsilon \equiv -\dot{H}/H^2$ is a measure of the degree to which the global symmetry of rigid time translations is broken. If $\epsilon \ll 1$ then a decoupling regime can exist near the Hubble scale.

The scales M_i and \bar{M}_i can be adjusted to reproduce the results of well-known models including canonical single-field inflation, Dirac–Born–Infeld inflation [17] and Ghost Inflation [18]. Alternatively they may be allowed to float. The action (2.4) then explores a range of interactions for fluctuations on a quasi-de Sitter background with nonlinearly realized Lorentz invariance, subject to the proviso (as described above) that only the dominant term for each M_i and \bar{M}_i has been retained. In principle these mass scales depend on time, but because we are taking the time-dependence of background quantities to be very weak we will treat them as constants.

3 Calculation of the bispectrum

In this paper our aim is to estimate the parameters M_i , \bar{M}_i by using observations to indirectly reconstruct the two- and three-point functions $\langle \pi\pi \rangle$ and $\langle \pi\pi\pi \rangle$. By itself, π is not an observable and neither are its correlations: the measurable quantity is the temperature fluctuation $\delta T/T$ as a function of angular position on the sky. Typically this is decomposed into harmonics, generating corresponding amplitudes $a_{\ell m}$,

$$\frac{\delta T(\hat{\mathbf{n}})}{T} = \sum_{\ell m} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}), \quad (3.1)$$

where $\hat{\mathbf{n}}$ represents an orientation on the sky and $Y_{\ell m}(\hat{\mathbf{n}})$ is a conventionally-normalized spherical harmonic. The amplitude $a_{\ell m}$ can be predicted in terms of primordial quantities using the formula⁵

$$a_{\ell m} = 4\pi(-i)^\ell \int \frac{d^3 k}{(2\pi)^3} \Delta_\ell(k) \zeta(\mathbf{k}) Y_{\ell m}(\hat{\mathbf{k}}), \quad (3.2)$$

where the ‘curvature perturbation’ $\zeta = \delta \ln a(\mathbf{x}, t)$ represents a fluctuation in the local scale factor $a(\mathbf{x}, t)$. It can be related to π via $\zeta = -H\pi$ up to terms which vanish in the limit $k/aH \rightarrow 0$, where k is the Fourier mode under consideration and aH is the comoving wavenumber associated with the Hubble length.

In writing (3.2) we have assumed that, for each relevant Fourier mode, $\zeta(\mathbf{k})$ attains a practically time-independent value by some time during the radiation era. The transfer function $\Delta_\ell(k)$ describes the subsequent process by which this time-independent seed perturbation is taken up by fluctuations in the primordial plasma and propagated to the surface of last

⁵We have absorbed a conventional factor of 3/5 into the normalization of the transfer function.

scattering, where it constitutes a temperature fluctuation δT . Under these circumstances, Eq. (3.2) shows that the n -point functions of the $a_{\ell m}$ can be linearly related to the n -point functions of $\zeta(\mathbf{k})$, and therefore $\pi(\mathbf{k})$, provided we evaluate the curvature perturbation in (3.2) at a time when the $O(k/aH)$ corrections in the relationship between π and ζ are negligible.

Correlation functions of ζ .—Therefore, we must estimate the correlation functions of ζ at the time they achieve their constant values. It is this requirement which makes the study of multiple-field models challenging [8, 19], because it is difficult to predict in advance when the time-independent epoch will occur. In single-field models the situation is simpler because the approximate global symmetry under rigid time translations (together with certain technical assumptions) is sufficient to prove the operator statement $\dot{\zeta} = 0$ in the limit $k/aH \rightarrow 0$ [20–25]. Therefore all correlation functions of ζ are constant on superhorizon scales, where $|k/aH|$ is negligible. An important consequence of this result is that subleading terms in the effective action (2.4) map to subleading terms in each n -point function [19], so to obtain a lowest-order result there is no need to consider corrections to (2.4) due to our neglect of time dependence in the M_i , \bar{M}_i .

In perturbation theory, a three- or higher n -point function is computed by integrating the reaction rate for an n -body interaction together with factors representing the available interaction volume and the probability for suitable particles to be present. These techniques were first applied to inflation by Maldacena [26] and later refined by various authors [17, 27–33]. We refer to this literature for technical details. In this section we wish to emphasize that, in the context of a general effective field theory, there are subtleties associated with computation of the field mode functions. These represent the amplitude for single-particle excitations of the vacuum. Therefore their properties significantly influence the n -point functions because they determine the probability for particles to be present in the interaction region.

Bartolo et al. noted that the scales M_1 , \bar{M}_1 , \bar{M}_2 and \bar{M}_3 in Eq. (2.4) are correlated with contributions to the second-order effective action, and of these \bar{M}_2 and \bar{M}_3 generate kinetic terms involving fourth-order derivatives. Kinetic terms of this type had previously been encountered in the ‘Ghost Inflation’ scenario proposed by Arkani-Hamed et al. [18]. Such terms are problematic because they imply that the mode functions can no longer be expressed in terms of elementary functions. This obstructs analytic integration of the interaction rate and hence each n -point function. In scenarios which require these high-order kinetic terms, exact results for the correlation functions typically require numerical calculation.

Bartolo et al. gave an explicit formula for the mode functions including the contribution of fourth-order terms, expressed in terms of hypergeometric functions and generalized Laguerre polynomials [13], and performed an analysis of its influence on each n -point function [13, 14]. They concluded that the fourth-order terms could significantly modify propagation deep within the horizon, but produced qualitatively similar results near the epoch of horizon exit. Since the correlations we are seeking to study are exponentially dominated by interactions occurring near this epoch, this implies that an acceptable estimate of the bispectrum *shape* can be obtained using a simpler mode function which does not account for fourth-order contributions. The penalty for this approximation is an uncertainty in the amplitude, which arises from a difference in *normalization* between the mode functions with and without the inclusion of fourth-order terms. For more details we refer to the discussion in Refs. [13, 14].

In this paper we follow Bartolo et al. and estimate each bispectrum shape by neglecting

the influence of fourth-order terms. This means that our results must be interpreted with some care:

1. When applied to a model for which $\bar{M}_2 = \bar{M}_3 = 0$, our results are exact within the approximations which have already been discussed. In this case, we expect both our qualitative and quantitative conclusions to be reliable.
2. When applied to a model for which at least one of \bar{M}_2 or \bar{M}_3 is nonzero, the normalization of our bispectra will be incorrect for the reasons just explained. This uncertainty in normalization affects the bispectrum for each operator in Eq. (2.4), not just those associated with the scales \bar{M}_2 and \bar{M}_3 —but we expect that it should be approximately the same for all of them. In this scenario, our quantitative estimates for the mass scales M_i , \bar{M}_i are not reliable. However, qualitative conclusions regarding the relative importance of each operator should be unaffected because ratios of these mass scales divide out any uncertainty in normalization.

To obtain reliable quantitative estimates of the mass scales when at least one of \bar{M}_2 or \bar{M}_3 is nonzero, it would be necessary to substitute numerical calculations of the bispectra in our analysis. In addition, the likelihood function to be discussed in §4 would no longer be approximately Gaussian and the analysis to follow should be replaced by a more sophisticated numerical exploration of the likelihood surface.

These modifications significantly increase the complexity of the analysis. They would certainly be required if observations provided pressure to include an \bar{M}_2 or \bar{M}_3 term in the effective Lagrangian. At present, our view is that such a step in complexity is not justified by the data.

4 Estimating the EFT mass scales

Bispectrum of curvature perturbation.—Under the approximations discussed in §3, the shapes of the bispectra generated by each operator in Eq. (2.4) were plotted in Ref. [13]. We tabulate analytical results for the corresponding three-point functions (which were not given explicitly in Ref. [13]) in Appendix A. The total three-point function for ζ should be obtained by summing these contributions, weighted by an appropriate mass scale M_i or \bar{M}_i .

In what follows it will be convenient to collect these mass scales, together with other normalization factors, into dimensionless combinations λ_α given in Table 1. There are eleven independent mass scales and therefore eleven independent λ_α . We use Greek indices α, β, \dots , to label these scales and the corresponding Lagrangian operators, which we write abstractly as \mathcal{O}^α . Each index ranges over the values A, B, \dots, K , and the effective action is the combination $S_{\text{EFT}} = \int d^4x \sqrt{-g} \sum_\alpha \lambda_\alpha \mathcal{O}^\alpha$. We position indices so that the normal rules of the Einstein summation convention are respected, but for clarity we will usually write summations over these indices explicitly. With these choices, we find

$$B_\zeta(k_1, k_2, k_3) = \frac{3}{5} \sum_\alpha \lambda_\alpha B^\alpha(k_1, k_2, k_3), \quad (4.1)$$

where B labels the bispectrum, defined so that (for example)

$$\langle \zeta(\mathbf{k}_1) \zeta(\mathbf{k}_2) \zeta(\mathbf{k}_3) \rangle = (2\pi)^3 \delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_\zeta(k_1, k_2, k_3), \quad (4.2)$$

Parameter	expressed as a mass scale	
	in terms of H	H eliminated
λ_A	$-\frac{65}{20736} \frac{1}{\pi^4 \epsilon^3 c_s^4 \mathcal{A}_s^2} \frac{\bar{M}_1^3 H^3}{M_P^6}$	$-\frac{65}{648\sqrt{2}} \frac{1}{\pi \epsilon^{3/2} c_s^{5/2} \mathcal{A}_s^{1/2}} \frac{\bar{M}_1^3}{M_P^3}$
λ_B	$-\frac{85}{10368} \frac{1}{\pi^4 \epsilon^3 c_s^2 \mathcal{A}_s^2} \frac{M_2^4 H^2}{M_P^6}$	$-\frac{85}{1296} \frac{1}{\pi^2 \epsilon^2 c_s \mathcal{A}_s} \frac{M_2^4}{M_P^4}$
λ_C	$-\frac{325}{62208} \frac{1}{\pi^4 \epsilon^3 c_s^4 \mathcal{A}_s^2} \frac{\bar{M}_2^2 H^4}{M_P^6}$	$-\frac{325}{972} \frac{1}{\epsilon c_s^2} \frac{\bar{M}_2^2}{M_P^2}$
λ_D	$\frac{5}{3888} \frac{1}{\pi^4 \epsilon^3 \mathcal{A}_s^2} \frac{M_3^4 H^2}{M_P^6}$	$\frac{5}{486} \frac{c_s}{\pi^2 \epsilon^2 \mathcal{A}_s} \frac{M_3^4}{M_P^4}$
λ_E	$-\frac{65}{7776} \frac{1}{\epsilon^3 c_s^4 \mathcal{A}_s^2} \frac{\bar{M}_3^2 H^4}{M_P^6}$	$-\frac{130}{243} \frac{1}{\epsilon c_s^2} \frac{\bar{M}_3^2}{M_P^2}$
λ_F	$\frac{5}{3888} \frac{1}{\pi^4 \epsilon^3 c_s^2 \mathcal{A}_s^2} \frac{\bar{M}_4^3 H^3}{M_P^6}$	$\frac{5\sqrt{2}}{243} \frac{1}{\pi \epsilon^{3/2} c_s^{1/2} \mathcal{A}_s^{1/2}} \frac{\bar{M}_4^3}{M_P^3}$
λ_G	$-\frac{65}{46656} \frac{1}{\pi^4 \epsilon^3 c_s^4 \mathcal{A}_s^2} \frac{\bar{M}_5^2 H^4}{M_P^6}$	$-\frac{65}{729} \frac{1}{\epsilon c_s^2} \frac{\bar{M}_5^2}{M_P^2}$
λ_H	$-\frac{65}{186624} \frac{1}{\pi^4 \epsilon^3 c_s^4 \mathcal{A}_s^2} \frac{\bar{M}_6^2 H^4}{M_P^6}$	$-\frac{65}{2916} \frac{1}{\epsilon c_s^2} \frac{\bar{M}_6^2}{M_P^2}$
λ_I	$\frac{115}{69984} \frac{\bar{M}_7 H^5}{\pi^4 \epsilon^3 c_s^6 \mathcal{A}_s^2} \frac{1}{M_P^6}$	$\frac{460\sqrt{2}}{2187} \frac{\pi \mathcal{A}_s^{1/2}}{\epsilon^{1/2} c_s^{7/2}} \frac{\bar{M}_7}{M_P}$
λ_J	$\frac{115}{279936} \frac{1}{\pi^4 \epsilon^3 c_s^6 \mathcal{A}_s^2} \frac{\bar{M}_8 H^5}{M_P^6}$	$\frac{115\sqrt{2}}{2187} \frac{\pi \mathcal{A}_s^{1/2}}{\epsilon^{1/2} c_s^{7/2}} \frac{\bar{M}_8}{M_P}$
λ_K	$-\frac{115}{559872} \frac{1}{\pi^4 \epsilon^3 c_s^6 \mathcal{A}_s^2} \frac{\bar{M}_9 H^5}{M_P^6}$	$-\frac{115}{2187\sqrt{2}} \frac{\pi \mathcal{A}_s^{1/2}}{\epsilon^{1/2} c_s^{7/2}} \frac{\bar{M}_9}{M_P}$

Table 1: Parameters λ_α in terms of the coefficients in the Lagrangian

and similarly for the π three-point function, which produces a bispectrum B^α for each operator \mathcal{O}^α . In Eq. (4.1) the normalization of each λ_α has been adjusted so that the B^α satisfy

$$\frac{B^\alpha(k, k, k)}{6P_\zeta(k)^2} = 1, \quad (4.3)$$

where $P_\zeta(k) = 2\pi^2 \mathcal{A}_s / k^3$ is the power spectrum and \mathcal{A}_s is the scalar amplitude. Each bispectrum is evaluated at the equilateral point and in principle depends on the side length k . However, because the n -point functions we study are nearly scale invariant (which for the bispectra implies $B^\alpha(k, k, k) \sim k^{-6}$), the precise choice of scale used to fix this normalization is unimportant. For a precisely local bispectrum, our convention (4.3) would make the corresponding λ_α equal to the conventional nonlinearity parameter $f_{\text{NL}}^{\text{local}}$. In general, however, the B^α will not be local and although each nonlinearity parameter such as $f_{\text{NL}}^{\text{local}}$, $f_{\text{NL}}^{\text{equi}}$, etc., will

be a linear combination of the λ_α , the coefficients in these combinations need not be simple.

Projection to the CMB bispectrum.—Eq. (3.2) shows that measurements of the microwave background anisotropies do not furnish information about B_ζ directly, but only via correlation functions of the $a_{\ell m}$. The first such correlation function which contains accessible information regarding B_ζ is the three-point function $\langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} \rangle$. It is conventional to extract a combinatorical factor $\mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3}$ —the so-called ‘Gaunt integral’—which is nonzero only for allowed combinations of the ℓ_i and m_i . The remainder of the correlation function is written as a ‘reduced bispectrum’ $b_{\ell_1 \ell_2 \ell_3}$,

$$\langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} \rangle = b_{\ell_1 \ell_2 \ell_3} \mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3}. \quad (4.4)$$

Our task is to determine $b_{\ell_1 \ell_2 \ell_3}$ given B_ζ . A strategy for doing so was developed by Fergusson, Shellard and collaborators [34–38] and extended by other authors [39, 40]. We briefly recount the steps in this strategy, using the notation of Refs. [41, 42]. First, for each bispectrum B^α one defines a corresponding dimensionless ‘shape function’ S^α using a fixed reference bispectrum B_{ref} ,

$$S^\alpha(k_1, k_2, k_3) \equiv \frac{B^\alpha(k_1, k_2, k_3)}{B_{\text{ref}}(k_1, k_2, k_3)}. \quad (4.5)$$

In principle, our final predictions do not depend on the choice of B_{ref} . In practice we will be forced to make approximations, some of which may introduce a residual dependence on B_{ref} . For this reason it is helpful to choose a form which has good numerical properties; often it is a good choice to fix a B_{ref} which shares features similar to the B^α . In this paper we will use the ‘constant’ bispectrum [43],

$$B_{\text{ref}}(k_1, k_2, k_3) = 6 \left(\frac{2\pi^2 \mathcal{A}_s}{k_1 k_2 k_3} \right)^2. \quad (4.6)$$

Second, one chooses a set of functions \mathcal{R}^n which furnish at least an approximate basis for the functions S^α . We define coefficients α_n^α so that

$$S^\alpha(k_1, k_2, k_3) \approx \sum_n \alpha_n^\alpha \mathcal{R}^n(k_1, k_2, k_3). \quad (4.7)$$

In practice it is only possible to retain a finite number of the \mathcal{R}^n ,⁶ so they should be chosen to give an acceptable approximation for each S^α using only a small number of modes. For details regarding the construction of suitable \mathcal{R}^n we refer to the literature [37, 39, 40]. It follows that, to a good approximation, the ζ bispectrum can be written

$$B_\zeta(k_1, k_2, k_3) \approx \frac{3}{5} B_{\text{ref}}(k_1, k_2, k_3) \sum_n \sum_\alpha \lambda_\alpha \alpha_n^\alpha \mathcal{R}^n(k_1, k_2, k_3). \quad (4.8)$$

The map from $\zeta(\mathbf{k})$ to $a_{\ell m}$ expressed by Eq. (3.2) is linear, and therefore the observable quantity $b_{\ell_1 \ell_2 \ell_3}$ must be proportional to a linear combination of the coefficients $\sum_\alpha \lambda_\alpha \alpha_n^\alpha$. Therefore we can write

$$b_{\ell_1 \ell_2 \ell_3} = \sum_{n,m} \Gamma_n^m b_{\ell_1 \ell_2 \ell_3}^n \sum_\alpha \lambda_\alpha \alpha_m^\alpha = \sum_\alpha \lambda_\alpha b_{\ell_1 \ell_2 \ell_3}^\alpha, \quad (4.9)$$

⁶The error associated with this truncation is one place where residual dependence on the reference bispectrum B_{ref} can appear.

where $b_{\ell_1\ell_2\ell_3}^\alpha$ is the reduced angular bispectrum associated with the operator \mathcal{O}^α ,

$$b_{\ell_1\ell_2\ell_3}^\alpha \equiv \sum_{n,m} \alpha_m^\alpha \Gamma_n^m b_{\ell_1\ell_2\ell_3}^n, \quad (4.10)$$

and the basis functions $b_{\ell_1\ell_2\ell_3}^n$ are defined in Refs. [34, 41]. They do not depend on any details of the cosmological model, which is carried only by the ‘transfer matrix’ Γ_n^m . This can be expressed as an integral over the linear transfer function $\Delta_\ell(k)$. The virtue of the approach of Fergusson, Shellard et al. is that calculation of Γ_n^m is numerically more tractable than calculation of an arbitrary bispectrum. Explicit formulae for the $b_{\ell_1\ell_2\ell_3}^n$ and Γ_n^m were given in Refs. [41, 42]. To compress notation we define $\bar{\alpha}_n^\alpha \equiv \sum_m \Gamma_n^m \alpha_m^\alpha$ and $\bar{\beta}_n = \sum_\alpha \lambda_\alpha \bar{\alpha}_n^\alpha$, from which it follows that

$$b_{\ell_1\ell_2\ell_3} \approx \sum_n \bar{\beta}_n b_{\ell_1\ell_2\ell_3}^n. \quad (4.11)$$

This projection procedure introduces correlations between the observable bispectra $b_{\ell_1\ell_2\ell_3}^\alpha$ produced by different Lagrangian operators, even if the corresponding primordial bispectra $B^\alpha(k_1, k_2, k_3)$ are nearly uncorrelated. We will return to this issue in §4.1 below.

Comparison with data.—It follows from Eq. (4.11) that information about the observable bispectrum from a microwave background survey can be reduced to estimates of the $\bar{\beta}_n$ and their covariances. We denote these estimates $\hat{\beta}_n$ and write their covariance matrix $\hat{\mathcal{C}}_{mn}$,

$$\hat{\mathcal{C}}_{mn} \approx \langle \Delta \hat{\beta}_m \Delta \hat{\beta}_n \rangle, \quad (4.12)$$

where $\Delta \hat{\beta}_n$ is the deviation of the observed $\hat{\beta}_n$ from its expected value, $\Delta \hat{\beta}_n \equiv \hat{\beta}_n - \bar{\beta}_n$. The standard methods of linear algebra can be used to obtain an orthonormal combination of bispectra from a Cholesky decomposition of this matrix [34–38, 41, 42]. In the interests of simplicity we assume this has been done, which makes $\hat{\mathcal{C}}_{mn}$ equal (for the rotated bispectra) to the identity matrix.⁷

For a set of measurements $\hat{\beta}_n$, the likelihood function \mathcal{L} represents the probability that these values would be observed given a particular model for their origin—in this case, the effective Lagrangian (2.4) with parameters λ_α . Assuming that the $\hat{\beta}_n$ are Gaussian distributed, this probability can be written

$$\mathcal{L}(\hat{\beta}_n | \lambda_\alpha) = \frac{1}{\sqrt{2\pi \det \hat{\mathcal{C}}}} \exp \left(-\frac{1}{2} \sum_{m,n} (\hat{\mathcal{C}}^{-1})^{mn} \Delta \hat{\beta}_m \Delta \hat{\beta}_n \right). \quad (4.13)$$

⁷We estimate $\hat{\mathcal{C}}_{mn}$ from the covariance matrix of the cubic needlet statistic, after changing basis to the \mathcal{R}_n as described in §II.B of Ref. [41]. (See also Ref. [42].) The signal-to-noise for the bispectrum (4.12) is roughly

$$\left(\frac{S}{N} \right)^2 \approx -2 \ln \mathcal{L} \approx \sum_{mn} \bar{\beta}_m (\hat{\mathcal{C}}^{-1})^{mn} \bar{\beta}_n,$$

making $\hat{\mathcal{C}}_{mn}$ a Fisher estimate of the covariance for the $\bar{\beta}_n$. Under the assumption that the bispectrum is small we assume that this covariance matrix is a reasonable approximation to $\langle \Delta \beta_m \Delta \beta_n \rangle$.

To compute $\hat{\mathcal{C}}_{mn}$ we use a suite of 50,000 Gaussian simulations, and for the change-of-basis coefficients we use a suite of 1,000 non-Gaussian simulations. These simulations incorporate the effect of the WMAP beam and mask for each channel, including noise with variance-per-pixel determined by the WMAP 9-year data release. In the rotated basis, where we choose $\hat{\mathcal{C}}_{mn} = \delta_{mn}$, all these details are transferred to the definition of the $\hat{\beta}_n$.

Maximum likelihood estimator.—It is now simple to construct a maximum likelihood estimator for the λ_α by finding the combination which has the greatest likelihood given the data. This gives the estimate

$$\hat{\lambda}_\alpha = \sum_\beta \hat{b}^\beta (\hat{\mathcal{F}}^{-1})_{\beta\alpha}, \quad (4.14)$$

where \hat{b}^α is defined by

$$\hat{b}^\alpha = \sum_{m,n} \hat{\beta}_m (\hat{\mathcal{C}}^{-1})^{mn} \bar{\alpha}_n^\alpha. \quad (4.15)$$

The matrix $\hat{\mathcal{F}}$ is the Fisher matrix associated with the likelihood (4.13),

$$\hat{\mathcal{F}}^{\alpha\beta} = -\frac{\partial^2 \ln \mathcal{L}}{\partial \lambda_\alpha \partial \lambda_\beta} = \sum_{m,n} \bar{\alpha}_m^\alpha (\hat{\mathcal{C}}^{-1})^{mn} \bar{\alpha}_n^\beta. \quad (4.16)$$

Truncating at the quadratic level, its inverse is formally the covariance matrix of the $\hat{\lambda}_\alpha$,

$$\langle (\hat{\lambda}_\alpha - \lambda_\alpha)(\hat{\lambda}_\beta - \lambda_\beta) \rangle = (\hat{\mathcal{F}}^{-1})_{\alpha\beta}. \quad (4.17)$$

In the rotated basis, where $\hat{\mathcal{C}}_{mn}$ is defined to be the unit matrix, Eq. (4.16) makes $\hat{\mathcal{F}}^{\alpha\beta}$ the square of the matrix $\bar{\alpha}_n^\alpha$ which expresses the decomposition of the angular bispectrum corresponding to the operator \mathcal{O}^α . In practice the Fisher formalism and Eq. (4.17) are likely to be trustworthy only in the limit of sufficiently high signal-to-noise.

The maximum likelihood estimator is an essentially frequentist concept, as is its variance (4.17). In a Bayesian framework one should instead interpret $(\hat{\mathcal{F}}^{-1})_{\alpha\beta}$ as the covariance of the posterior probability distribution of the parameters λ_α , constructed from a single set of measurements $\hat{\beta}_n$, assuming that any prior probabilities for the λ_α are flat over the range of interest.

4.1 How many independent shapes?

This analysis applies provided the matrix $\hat{\mathcal{F}}^{\alpha\beta}$ is invertible. However, invertibility may fail if two linear combinations of the operators \mathcal{O}^α produce nearly degenerate angular bispectra. This would imply that $\hat{\mathcal{F}}^{\alpha\beta}$ has an approximate null eigenvector.

The appearance of exact or approximate null eigenvectors implies that the likelihood function is a singular Gaussian distribution: it does not vary along directions in parameter space which correspond to the null eigenvectors. Therefore the variance of the maximum likelihood estimator (4.17) is formally infinite for all $\hat{\lambda}_\alpha$. To deal with this one should first discard those combinations of parameters which are unconstrained by the likelihood function. This is necessary even in the case of an approximate null eigenvector, because although the Fisher matrix may be formally invertible it will usually be ill-conditioned. Therefore we should trust a numerical inversion only if it is possible to compute $\hat{\mathcal{F}}^{\alpha\beta}$ to very high accuracy. Typically this cannot be done because the accuracy with which we know $\hat{\mathcal{F}}^{\alpha\beta}$ is limited by our ability to estimate $\hat{\mathcal{C}}_{mn}$, and by the numerical integrations required to compute $\bar{\alpha}_n^\alpha$. For a discussion of the issues involved in handling singular Fisher matrices, see (for example) Ref. [44].

Measures of correlation.—Two operators will produce degenerate bispectra if their decomposition coefficients α_m^α or $\bar{\alpha}_m^\alpha$ are nearly the same. These two measures do not have to agree, because (as explained on p. 11) the projection from α_m^α to $\bar{\alpha}_m^\alpha$ can change the degree

of correlation. The Fisher matrix $\hat{\mathcal{F}}$ is constructed from angular bispectra, and therefore—as a point of principle—the problematic degeneracies are those which occur for the $\bar{\alpha}_m^\alpha$. But in practice, for computation reasons, it is sometimes more practical to use the α_m^α as a proxy.

To measure the correlation between two primordial bispectra B_1 and B_2 we introduce an inner product, defined by

$$\langle\langle S^1, S^2 \rangle\rangle \equiv \int_{\mathcal{V}} dv S_1(k_1, k_2, k_3) S_2(k_1, k_2, k_3) \omega(k_1, k_2, k_3), \quad (4.18)$$

where S_1 and S_2 are the corresponding shape functions, dv is an element of volume on the integration domain \mathcal{V} (which corresponds to allowable triangular configurations of the momenta \mathbf{k}_i), and ω is a weight function which can be chosen to suit our convenience. For a more detailed discussion of Eq. (4.18) we refer to the literature [34]. We normalize the \mathcal{R}^n so that $\langle\langle \mathcal{R}^m, \mathcal{R}^n \rangle\rangle = \delta^{mn}$ and therefore (4.7) implies

$$\langle\langle S^1, S^2 \rangle\rangle = \sum_n \alpha_m^1 \alpha_m^2. \quad (4.19)$$

When measuring correlations between angular bispectra it is helpful to account for the ability of the WMAP instrument to distinguish between different shapes. This ability is measured by the matrix $(\hat{\mathcal{C}}^{-1})^{mn}$ discussed in footnote 7 on p. 11. We define

$$\langle\langle b_{\ell_1, \ell_2, \ell_3}^1, b_{\ell_1, \ell_2, \ell_3}^2 \rangle\rangle = \sum_{m, n} \bar{\alpha}_m^1 (\hat{\mathcal{C}}^{-1})^{mn} \bar{\alpha}_n^2. \quad (4.20)$$

Note that we write the inner product $\langle\langle \cdot, \cdot \rangle\rangle$ for both the primordial and angular bispectra, but they are *not* equal; the definition is different depending whether it is taken between primordial or angular bispectra. In either case it is conventional to measure the correlation between shapes by defining a ‘cosine’,

$$\cos(1, 2) = \frac{\langle\langle 1, 2 \rangle\rangle}{\langle\langle 1 \rangle\rangle^{1/2} \langle\langle 2 \rangle\rangle^{1/2}}, \quad (4.21)$$

where ‘1’ and ‘2’ should be substituted by the appropriate angular or primordial bispectrum.

Principal directions.—To factor out the degenerate directions we diagonalize $\hat{\mathcal{F}}$, finding a new orthogonal matrix \mathbf{U} and a nonnegative-definite diagonal matrix $\mathbf{\Sigma}$ so that $\hat{\mathcal{F}} = \mathbf{U} \mathbf{\Sigma} \mathbf{U}^t$ where a superscript ‘t’ denotes matrix transposition.⁸ The matrix \mathbf{U} can be regarded as a rotation from the operators \mathcal{O}^α to a new set of operators $\mathcal{O}^{\alpha'}$ which satisfy $\mathcal{O}^{\alpha'} = \sum_\alpha \mathcal{O}^\alpha U_\alpha^{\alpha'}$, and likewise a new set of dimensionless coefficients $\lambda_{\alpha'} = \sum_\alpha \lambda_\alpha U_\alpha^{\alpha'}$. The Lagrangian $\sum_\alpha \lambda_\alpha \mathcal{O}^\alpha = \sum_{\alpha'} \lambda_{\alpha'} \mathcal{O}^{\alpha'}$ is invariant under a rigid rotation of this kind.

The presence of degeneracies means that the eigenvalues of $\hat{\mathcal{F}}$ vary significantly in magnitude. The largest eigenvalues are

$$2.23 \times 10^{-3}, 1.70 \times 10^{-4}, 7.17 \times 10^{-7}, 1.22 \times 10^{-9}, \text{ and } 1.20 \times 10^{-14}, \quad (4.22)$$

⁸In practice, it can happen that numerical inaccuracies cause $\hat{\mathcal{F}}$ to develop very small negative eigenvalues which spoil simple diagonalization strategies. Where this occurs we perform a singular value decomposition, which corresponds to finding (possibly complex) unitary matrices \mathbf{U} , \mathbf{V} and a nonnegative-definite diagonal matrix $\mathbf{\Sigma}$ so that $\hat{\mathcal{F}} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^t$. We discard complex directions and check that the results are stable under exchange of \mathbf{U} and \mathbf{V} .

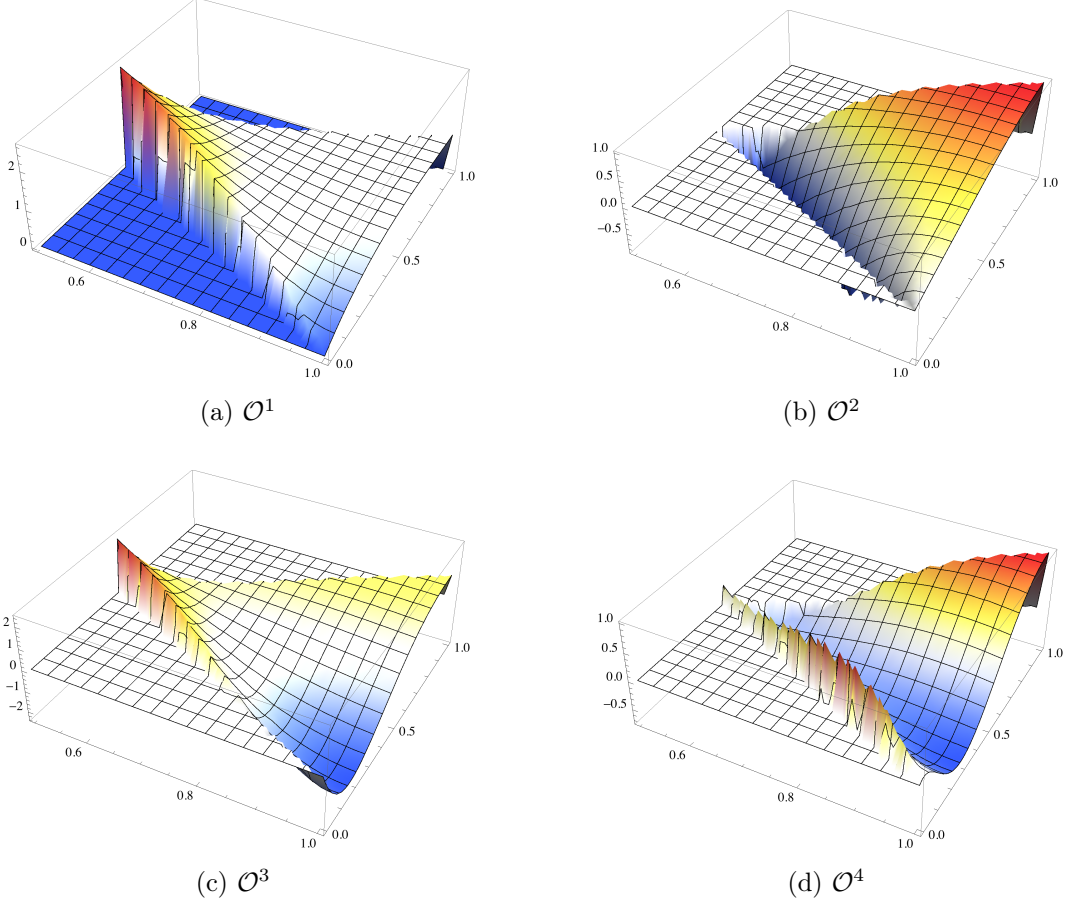


Figure 1: Bispectrum shapes generated by the operators $\mathcal{O}^{\alpha'}$ corresponding to the constrainable parameter combinations $\lambda_{\alpha'}$. These plots follow the conventions of Babich et al. [45]. For each $\mathcal{O}^{\alpha'}$ the plotted quantity is $x^2 y^2 B_{\alpha'}(x, y, 1) / B_{\alpha'}(1, 1, 1)$ where $x = k_1/k_3$ and $y = k_2/k_3$ (no sum on α').

with the remaining eigenvalues being of order 10^{-15} or smaller. We retain the first four, which corresponds to a hierarchy between largest and smallest eigenvalues of $\sim 10^6$. The corresponding eigenvectors in parameter space yield four linear combinations $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ which can be constrained. It should be remembered that because the covariance matrix $\hat{\mathcal{C}}_{mn}$ defined in (4.12) depends on details of the WMAP experiment (including the masks, beam and noise properties as described in footnote 7 on p. 11), the Fisher matrix and therefore the shapes corresponding to these leading eigenvalues also depend on these details. They may vary between experiments, depending on the varying sensitivity of each experiment to different regions of multipole-space. The precise specification of the leading shapes given in Table 3, and we plot the shapes of the corresponding primordial bispectra in Fig. 1.

Shapes of principal directions.—These shapes can be given an approximate interpretation in terms of the standard templates. Fig. 1a is associated with the largest eigenvalue, and is therefore the best-measured shape. It exhibits significant correlations for $x = y = 0.5$, which corresponds the ‘folded’ configuration [46]. Fig. 1b exhibits significant correlations in the equilateral limit $x = y = 1$, and some anticorrelation in the folded configuration. It can be

regarded as an approximate ‘orthogonal’ shape [12]. Together, a linear combination of these two configurations can be used to produce an approximate ‘equilateral’ shape. These results are consistent with the forecast of Byun & Bean [39], who suggested that (neglecting the local shape), the highest signal-to-noise should be achieved for shapes similar to the folded and orthogonal templates. Note that Byun & Bean’s analysis was based on a survey with Planck-like masks, beams and noise rather than the WMAP9 characteristics adopted here.

Fig. 1c has an interior node, where—without our choice of signs—anti-correlations have a local maximum in the interior of the allowed triangular region. This is quite different to the behaviour of Figs. 1a–1b, in which local maxima only occur for extreme configurations on the boundary of the allowed region. The shape of Fig. 1c is similar to a shape produced in a Galileon theory by Creminelli et al. [47], and later reproduced in a general Horndeski Lagrangian by Refs. [32, 48]. Finally, Fig. 1d is a complex shape containing an interior node together with substantial correlations in the equilateral configuration. It represents something different from the shape of Fig. 1c, but it will be seen in §4.2 below that it is rather weakly constrained by the data.

These results are consistent with the conclusions of Ribeiro et al. [48], who found that in a very general single-field model⁹ it could be possible to produce a measurable signal in a mode similar to that of Fig. 1c, or equivalently the Creminelli et al. shape [47], but that further orthogonal shapes would be difficult to measure.

Correlation of shapes.—We tabulate the correlation between these shapes in Table 2, and also between these shapes and the standard CMB templates. The correlation is computed for the primordial bispectra using (4.18). In particular we note that, although the angular bispectra for the $\mathcal{O}^{\alpha'}$ are orthogonal by construction, mapping back to the primordial bispectrum introduces some correlation; for example, $\cos(\mathcal{O}^3, \mathcal{O}^4) = -0.62$. This degradation is expected, because the increasing covariance represented by (4.22) will cause noise to dominate over signal. Therefore the linear relationship between the primordial and CMB n -point functions, implied by Eq. (3.2), is no longer satisfied.

One can regard these results as a reflection of the fact that the first three operators \mathcal{O}^1 , \mathcal{O}^2 and \mathcal{O}^3 are reasonably well-measured, whereas the fourth operator \mathcal{O}^4 is only weakly constrained.

4.2 Results

A framework for estimating the $\hat{\beta}_n$ from a CMB temperature map using wavelet or needlet methods was developed by Regan et al. [41, 42]. We apply these methods to 9-year data from the WMAP satellite [2, 49]. For the constrainable parameters $\{\lambda_{\alpha'}\} = \{\lambda_1, \lambda_2, \lambda_3, \lambda_4\}$ we find

	Estimate
$\hat{\lambda}_1$	-22.9 ± 20.9
$\hat{\lambda}_2$	94.9 ± 76.7
$\hat{\lambda}_3$	-956 ± 1180
$\hat{\lambda}_4$	42400 ± 28600

The quoted errors are 1σ and marginalized over the other $\lambda_{\alpha'}$.

⁹Ribeiro et al. worked with a model for the fluctuations which is equivalent to the fluctuations in a general Horndeski action. Although very permissive, this model is still less general than the full effective field theory (2.4).

	\mathcal{O}^1	\mathcal{O}^2	\mathcal{O}^3	\mathcal{O}^4
\mathcal{O}^1	1.00	-0.03	-0.11	-0.01
\mathcal{O}^2	-0.03	1.00	0.17	0.24
\mathcal{O}^3	-0.11	0.17	1.00	-0.62
\mathcal{O}^4	-0.01	0.24	-0.62	1.00
constant	-0.95	-0.21	-0.09	0.00
equilateral	-0.80	-0.57	0.03	-0.16
flat	-0.73	0.19	-0.38	0.31
local	-0.54	0.00	-0.29	0.02
orthogonal	0.36	-0.79	0.27	-0.35

Table 2: Cosines of the shapes appearing in Fig. 1 between themselves and the standard CMB templates. Inner products are computed using (4.18) and the constant bispectrum as a reference.

Senatore, Smith & Zaldarriaga [12] obtained constraints on the amplitude of the ‘equilateral’ and ‘orthogonal’ bispectrum templates from the 5-year WMAP data, and used these to constrain a subset of terms in the effective Lagrangian (2.4). They concluded that each shape included in their analysis could be approximately described by a linear combination of these two templates, up to $\sim 90\%$ correlation. However, they included only two of the operators in Eq. (2.4). Our analysis demonstrates that it is possible to increase the number of linearly independent operators from two to four, although the estimate $\hat{\lambda}_4 = 42400 \pm 28600$ shows that that sensitivity is already decreasing markedly for the fourth parameter.

5 Constraints on models

In §4 we obtained constraints on certain linear combinations of the EFT scales M_i , \bar{M}_i . The remaining linear combinations formally have infinite uncertainties because of degeneracies. Together, these results summarize the information which can be recovered from the WMAP9 bispectrum, but to apply them to specific models we must first match the mass scales M_i , \bar{M}_i . In this section we give two examples of this programme for models of observational interest: the Dirac–Born–Infeld model (‘DBI inflation’) and ‘Ghost inflation’.

Once the M_i , \bar{M}_i are known, the results of §4 would give four constraints on different combinations of these scales. Depending how many scales are needed to parametrize an individual model, it may be possible to estimate some or all of them, or they may even be over-constrained. The latter possibility indicates that the model is a poor fit for the data. Where more than four mass scales are needed to characterize a model, the constraints pick out an observationally-allowed subspace which is consistent with the CMB bispectrum measurements.

Methodology.— To map our four constraints for $\{\lambda_1, \lambda_2, \lambda_3, \lambda_4\}$ onto a subset of the original parameter space M_i , \bar{M}_i we minimize the value

$$X^2 = \sum_{\alpha'=1}^4 \frac{[\lambda_{\alpha'}(M_i, \bar{M}_i) - \hat{\lambda}_{\alpha'}]^2}{\Sigma_{\alpha'\alpha'}^{-1}}, \quad (5.1)$$

where $\Sigma_{\alpha'\beta'}$ is the diagonal matrix of principal eigenvalues listed in Eq. (4.22); $\lambda_{\alpha'}(M_i, \bar{M}_i)$ represents the value of the linear combination $\lambda_{\alpha'}$ which would be predicted given a fixed choice of mass scales M_i, \bar{M}_i ; and $\hat{\lambda}_{\alpha'}$ represents the value estimated from the data in §4. With four constraints on the $\lambda_{\alpha'}$ we can constrain up to four of the M_i, \bar{M}_i .

We take X^2 to be χ^2 -distributed with four degrees of freedom. The $\lambda_{\alpha'}$ are constructed from a linear combination of the $\hat{\beta}_n$, and we assume that the experimental error for each $\hat{\beta}_n$ is independent and Gaussian-distributed. Because the $\lambda_{\alpha'}$ are chosen to be orthogonal, the experimental errors on each $\hat{\lambda}_{\alpha'}$ will be obtained from a nearly uncorrelated sum of Gaussians, and will therefore also be nearly independent. This makes X^2 approximately equal to a sum of four approximately independent, unit Gaussians, and hence roughly χ^2 -distributed.

Confidence intervals for the M_i, \bar{M}_i could be determined by searching for suitable critical values of the χ^2 distribution. Alternatively, assuming that the $\hat{\lambda}_{\alpha'}$ have uncorrelated Gaussian errors, we could expand X^2 to second order around the maximum likelihood point,

$$\begin{aligned} X^2 &= X^2|_{\text{MLE}} + \sum_{\alpha'\beta'} \left. \frac{\partial^2 X^2}{\partial \lambda_{\alpha'} \partial \lambda_{\beta'}} \right|_{\text{MLE}} (\lambda_{\alpha'} - \lambda_{\alpha'}|_{\text{MLE}})(\lambda_{\beta'} - \lambda_{\beta'}|_{\text{MLE}}) + \cdots \\ &= X^2|_{\text{MLE}} + \Delta X^2, \end{aligned} \quad (5.2)$$

and construct confidence contours at the $n^{\text{th}}\text{-}\sigma$ level by searching for critical values where $\Delta X^2 = n^2$. In principle these methods agree if the $\hat{\lambda}_{\alpha'}$ are Gaussian and uncorrelated. We find that the agreement is not quite exact, which we ascribe to a small residual correlation between the errors on the $\hat{\lambda}_{\alpha'}$. The single-parameter constraints reported below are obtained using the second-order expansion (5.2), which reproduces the Fisher-matrix estimates. For two or more parameters we report constraints extracted from critical values of the full χ^2 -distribution with 4 degrees of freedom.

DBI inflation.—The first example we consider is the ‘Dirac–Born–Infeld’ or ‘DBI’ inflationary model.

The Dirac–Born–Infeld action describes fluctuations of a membrane moving in a warped transverse space, or ‘throat’. Under certain circumstances it can describe an inflationary epoch in which inflaton perturbations propagate at less than the speed of light from the perspective of a brane-based observer, due to constraints imposed by the extradimensional covering theory. The small sound speed means that these models can produce significant nongaussianities in the equilateral mode.

Fluctuations in a single-field DBI model can be described by the effective action (2.4), retaining only the B and D operators,

$$\lambda_B \mathcal{O}^B \propto M_2^4 \frac{1}{a^2} \dot{\pi} (\partial\pi)^2 \quad (5.3a)$$

$$\lambda_D \mathcal{O}^D \propto M_3^4 \dot{\pi}^3. \quad (5.3b)$$

This model does not involve the problematic scales \bar{M}_2, \bar{M}_3 which lead to normalization inaccuracies for the single-particle mode functions and therefore we expect our estimates to be quantitatively reliable.

The original DBI model had a single free parameter and therefore M_2 and M_3 cannot be chosen independently but are correlated as described below. Alternatively, one can consider a larger family of DBI-like models which retain only these operators but allow M_2 and M_3 to

vary. Following Senatore et al., constraints are typically expressed using the parameters

$$\frac{1}{c_s^2} = 1 - \frac{2M_2^4}{M_P^2 \dot{H}} = 1 - \frac{324}{85} \lambda_B, \quad (5.4a)$$

$$\tilde{c}_3 \left(\frac{1}{c_s^2} - 1 \right) = \frac{2M_3^4 c_s^2}{M_P^2 \dot{H}} = -\frac{243}{10} \lambda_D. \quad (5.4b)$$

Causality requires the speed of sound c_s to be less than unity. Since $\dot{H} < 0$ during inflation it follows that M_2^4 must be positive (see footnote 4 on p. 5). If $M_2^4 \gtrsim M_P^4 |\dot{H}|$ then a significant bispectrum can be generated. The reason for expressing constraints in terms of these parameters is that it is not possible to determine M_2 and M_3 without simultaneously specifying \dot{H} . The original DBI model imposes the constraint $\tilde{c}_3 = 3(1 - c_s^2)/2$.

We estimate the *joint* constraints on λ_B and λ_D to be

$$\lambda_B = -1151 \pm 760 \quad (5.5a)$$

$$\lambda_D = 946 \pm 584. \quad (5.5b)$$

The Planck collaboration expressed their constraints in terms of f_{NL} -like parameters $f_{\text{NL}}^{\text{EFT1}}$ and $f_{\text{NL}}^{\text{EFT2}}$. In our notation these correspond, respectively, to λ_B under the assumption $\lambda_D = 0$ and λ_D under the assumption $\lambda_B = 0$. Using the 2013 dataset, the Planck collaboration reported the bounds $f_{\text{NL}}^{\text{EFT1}} = 8 \pm 73$ and $f_{\text{NL}}^{\text{EFT2}} = 19 \pm 57$ [1]. Using the same notation, we find

$$f_{\text{NL}}^{\text{EFT1}} = 68.3 \pm 103 \quad (5.6a)$$

$$f_{\text{NL}}^{\text{EFT2}} = 69.3 \pm 79. \quad (5.6b)$$

The Planck2013 errors represent an improvement of order 30%.

Alternatively, each bound can be expressed in terms of c_s and \tilde{c}_3 . To compare with the constraints reported by the Planck collaboration we consider three possibilities. First, marginalizing over \tilde{c}_3 gives a conservative lower bound on c_s ,

$$c_s \geq 0.010 \quad \text{at 95\% confidence.} \quad (5.7a)$$

For comparison, Planck2013 found $c_s \geq 0.02$ [1] at the same confidence level. Second, imposing $\tilde{c}_3 = 0$ gives

$$c_s \geq 0.044 \quad \text{at 95\% confidence.} \quad (5.7b)$$

Finally, assuming the strict DBI relation between c_s and \tilde{c}_3 leaves c_s as a single free parameter. We find

$$c_s \geq 0.051 \quad \text{at 95\% confidence.} \quad (5.7c)$$

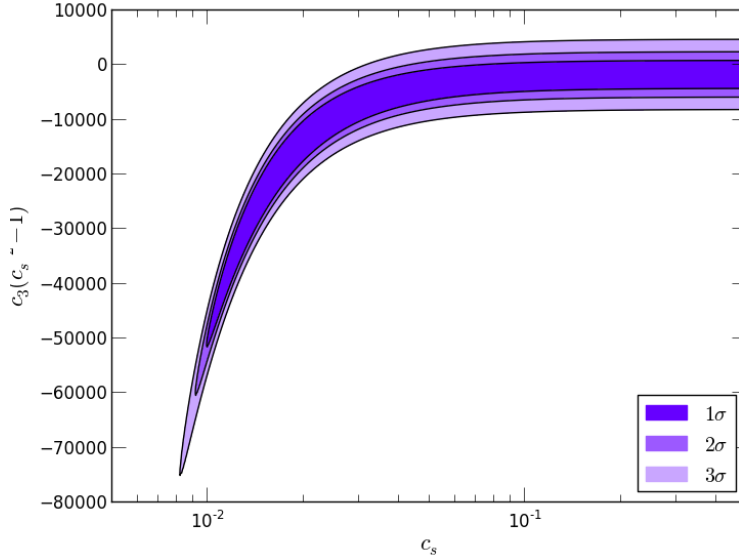
Planck2013 obtained $c_s \geq 0.07$ [1], also at 95% confidence. Eq. (5.7c) can also be expressed as an f_{NL} parameter for the DBI shape. This gives

$$f_{\text{NL}}^{\text{DBI}} = 69.6 \pm 97.4. \quad (5.8)$$

Finally, allowing both c_s and \tilde{c}_3 to vary results in a lower bound for c_s and relatively weak constraints for \tilde{c}_3 , plotted in Fig. 2.

Similar bounds were reported by Senatore et al. [12]. Our construction ensures that the bounds reported above correspond to the most accurate constraints which can be achieved

Figure 2: Constraints on the DBI-like parameters c_s , \tilde{c}_3 .



using this data set, because the shapes are explored using four rather than two orthogonal directions in the likelihood (5.1). For example, using only the leading principal component to construct the likelihood yields the constraint $c_s > 0.061$ in the DBI model. This bound unduly weights the component of the DBI shape along this principal direction, giving an overly optimistic constraint when compared with the four-component result (5.7c).

Ghost inflation.—Our second example is the ‘Ghost inflation’ model proposed by Arkani-Hamed et al. [18], in which inflation is driven by a so-called ‘ghost condensate’ which spontaneously breaks Lorentz invariance in the background. The ghost condensate is described by a scalar field ϕ whose time derivative gains a nonvanishing vacuum expectation value, $\langle \dot{\phi}^2 \rangle = M^2 \neq 0$. This expectation value is time-independent and is not diluted as inflation proceeds.

In the effective theory, fluctuations around the background correspond to nonzero \bar{M}_2^2 and \bar{M}_3^2 , and the limit $\dot{H} \rightarrow 0$. This limit sets the quadratic spatial-derivative terms in (2.4) to zero, so that the speed of sound is formally zero. The fluctuations are nevertheless propagating modes because higher-order spatial derivative terms are present in the Lagrangian. The relevant EFT operators are

$$\lambda_C \mathcal{O}^C \propto \frac{\bar{M}_2^4}{a^5} \left(\frac{H}{2} \partial^2 \pi (\partial \pi)^2 + \dot{\pi} \partial^2 \partial_i \pi \partial_i \pi \right) \quad (5.9a)$$

$$\lambda_E \mathcal{O}^E \propto \frac{\bar{M}_3^2}{a^4} \left(H \partial^2 \pi (\partial \pi)^2 + \dot{\pi} \partial^2 \partial_i \pi \partial_i \pi \right). \quad (5.9b)$$

The arrangement of derivatives is identical up to a relative factor of 2 in the first term. In

terms of the mass scales \bar{M}_2 and \bar{M}_3 we have

$$\lambda_C = -\frac{325}{972c_s^2\epsilon} \frac{\bar{M}_2^2}{M_{\text{P}}^2}, \quad (5.10a)$$

$$\lambda_E = -\frac{130}{243c_s^2\epsilon} \frac{\bar{M}_3^2}{M_{\text{P}}^2}. \quad (5.10b)$$

The inclusion of factors of c_s and ϵ is purely formal, since this model technically involves the limits $c_s \rightarrow 0$ and $\epsilon \rightarrow 0$. Proceeding as for the DBI model we obtain the maximum-likelihood estimates

$$\lambda_C = -3680 \pm 2280 \quad (5.11a)$$

$$\lambda_E = 3900 \pm 2450. \quad (5.11b)$$

Bartolo et al. observed that the operators (5.10a)–(5.10b) are nearly the same and chose to aggregate them into a single term operator by introducing a common mass scale \bar{M}_0 , satisfying by $\bar{M}_0^2 \equiv 2\bar{M}_3^2/3 \equiv -2\bar{M}_2^2$. The λ corresponding to this aggregate operator [still defined to satisfy the normalization condition (4.3)] represents an estimate of the amplitude of the ghost-inflation bispectrum, and we label it λ^{ghost} . As explained in §3, the ghost inflation model involves fourth-order kinetic terms whose details we do not capture, and therefore the precise normalization of this estimate is uncertain. We find

$$\lambda^{\text{ghost}} = -68.4 \pm 100.5. \quad (5.12)$$

For comparison, the Planck collaboration reported the constraint $f_{\text{NL}}^{\text{ghost}} = -23 \pm 88$. Both estimates agree that the bispectrum in this channel is consistent with zero within 1σ . In addition, this comparison shows that, even in a case where the normalization uncertainty is important, our result matches an exact calculation within a factor of order unity.

6 Model comparison using the bispectrum

The analyses of §§4–5 determine best-fit values for the λ_α , essentially in a frequentist sense, assuming a fixed model for the underlying microphysical fluctuations. Therefore our conclusions up to this point are restricted to *parameter estimation*.

Within this framework it is not possible to address questions such as whether the best-fit combination for DBI inflation, Eqs. (5.5a)–(5.5b), represents a better description of the data than the best-fit combination for Ghost inflation, Eqs. (5.11a)–(5.11b). These broader questions constitute the province of *model comparison*. Recent work has addressed the issue of inflationary model comparison based on measurements of the two-point function of the temperature anisotropy [50]. It is much more challenging to perform a similar analysis based on the three-point function. In this section we take some steps towards this objective within the framework described in §§2–4.

For computational reasons we must impose limitations on the meaning of the term ‘model’. Conceptually this should include whatever information is necessary to specify the value of each observable. For example, the transfer matrix Γ_n^m depends on the post-inflationary cosmological history and in a global analysis the parameters which specify this history should be varied in addition to the inflationary parameters M_i , \bar{M}_i . However, this generates a large parameter space which is expensive to search because calculation of Γ_n^m is

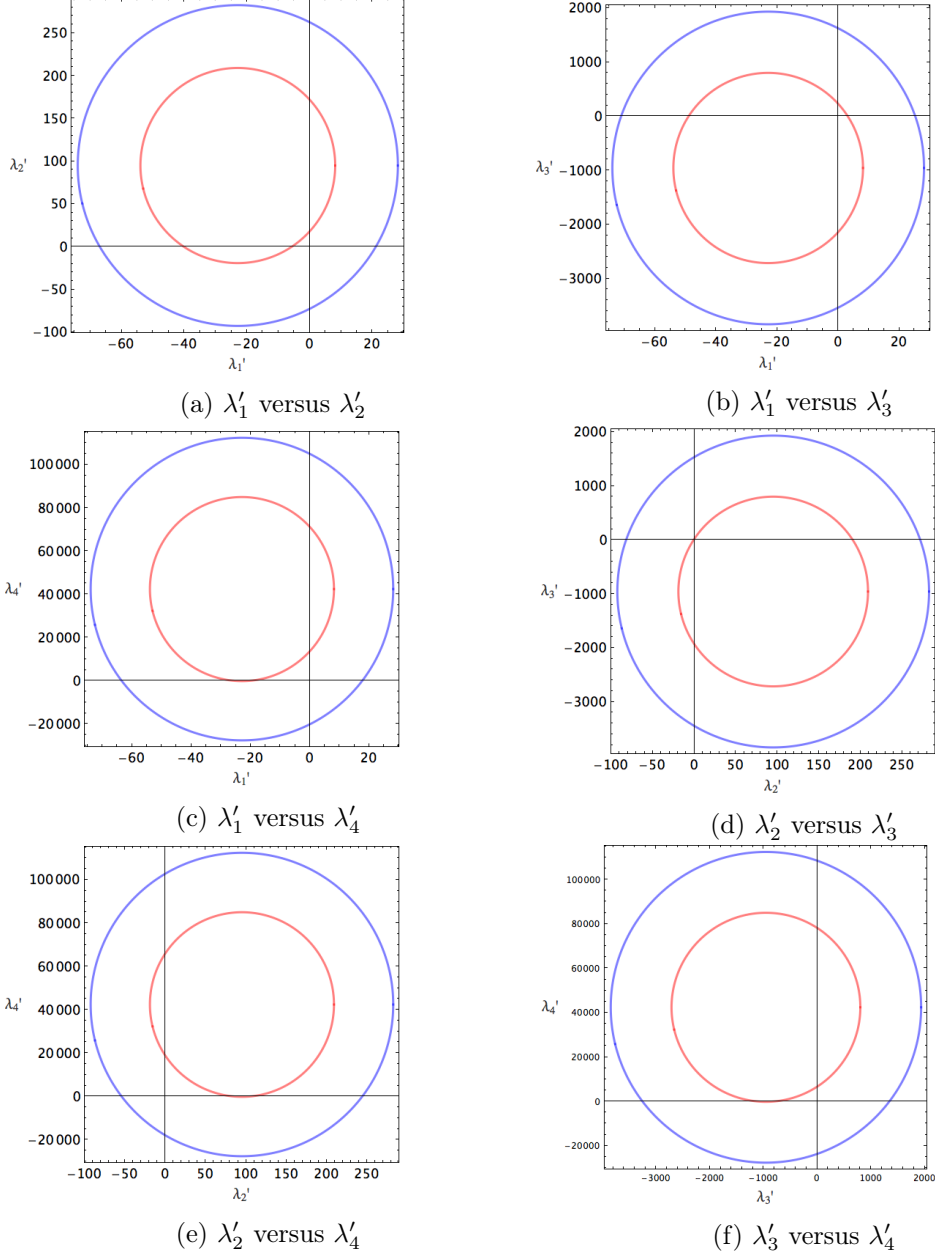


Figure 3: 1σ (red) 2σ (blue) confidence regions for two out of four principal components $\hat{\lambda}_{\alpha'}$ constrained by 9-year WMAP data. The results show consistency with zero magnitude generally within $1 - 1.5\sigma$, suggesting no strong evidence of nongaussianity in the single field inflationary parameter space.

time-consuming. In this analysis we will fix the transfer matrix using standard best-fit values for the post-inflationary history and address the more restricted question of which inflationary model yields a better fit for the three-point function given these assumptions.

Evidence for a model.—There is no single metric which unambiguously quantifies the evidence for or against a particular model. One choice is the ‘Bayes factor’. For a particular set of observations D and a pair of models M_1 and M_2 , this is defined to be the ratio of

likelihoods,

$$K_{12} = \frac{P(D|M_1)}{P(D|M_2)} = \frac{\int P(D|\lambda_1, M_1)P(\lambda_1|M_1) d\lambda_1}{\int P(D|\lambda_2, M_2)P(\lambda_2|M_2) d\lambda_2}. \quad (6.1)$$

We use λ_1, λ_2 to schematically denote two different choices of the parameters λ_α which characterize a particular model. Because M_1 and M_2 are different they may require a different number of parameters.

The probabilities $P(\lambda_i|M_i)$ represent the prior probability, for each model, that a particular parameter choice occurs. They must be chosen by hand. Where meaningful prior information exists (for example, previous measurements of a parameter) this can be encoded using these probabilities. But their arbitrariness implies that—unless it happens that K is nearly independent of the $P(\lambda|M)$ —the Bayes factor is not easy to interpret. Usually, K is independent of the priors only when the data are very constraining. In what follows we will see that the 9-year WMAP bispectrum data are insufficiently constraining for this to occur, so that ambiguities in the interpretation of K remain.

Empirical scales are used to give meaning to the Bayes factor. Commonly used examples are due to Jeffreys or Kass & Raftery [51]. In Kass & Raftery’s prescription, $\ln K$ in the range (1, 3) is considered evidence in favour of M_1 , whereas $\ln K$ in the range (3, 5) is considered strong evidence and larger values of K are considered decisive. Ratios for which $|\ln K| < 1$ are uninformative.

Choice of priors.—In our case the λ_α represent Lagrangian coefficients. Some prior estimates exist, but the datasets from which these were obtained are not independent of the 9-year WMAP data used in this analysis. For this reason we disregard these prior constraints, and therefore some other way must be found to justify the functional form of each prior.

If we insist that perturbation theory is valid then the λ_α should not be too large. This requirement is convenient but not obviously necessary. However, for the purpose of performing a concrete calculation we shall adopt it in what follows. In that case, the requirement that the bispectrum generated by the operator \mathcal{O}^α does not overwhelm the power spectrum \mathcal{P}_ζ is roughly $|\lambda_\alpha|\mathcal{P}_\zeta^{1/2} \lesssim 1$, and therefore $|\lambda_\alpha| \lesssim 10^4$. This limit is helpful but gives no guidance regarding the functional form of $P(\lambda_\alpha|M)$. To explore the range of outcomes we consider two possibilities:

- The ‘Jeffries prior’ $P(\lambda_\alpha) \propto |\lambda_\alpha|^{-1}$. This choice assigns equal probability for each decade of $|\lambda_\alpha|$: that is, for λ_α to be between 1 and 10, 10 and 100, and so on. The Jeffries prior makes it relatively likely for $|\lambda_\alpha|$ to be near zero, and therefore can be regarded as conservative.¹⁰
- The flat prior, for which $P(\lambda_\alpha)$ is constant. This choice assigns equal probability to each value of λ_α , and therefore makes it relatively more likely for $|\lambda_\alpha|$ to be large. It is less conservative than the Jeffries prior in the sense that that it enhances the probability for the Lagrangian (2.4) to predict observably large nongaussianities.

¹⁰Strictly, the Jeffries prior has a divergence at $\lambda_\alpha = 0$. We regularize this by cutting out the region $|\lambda_\alpha| < 1$ and taking $P(\lambda_\alpha)$ to be zero within it. We have checked that our results are robust to modest changes of the boundary value. The choice of cutoff at unity is, of course, somewhat arbitrary. However, motivated by the fact that in the single parameter case λ_α corresponds to the conventionally defined f_{NL} parameter, we note that error bars on f_{NL} for single field inflationary models are at best expected to achieve values of order unity. Therefore, we shall regard $\lambda_\alpha = 1$ as a ‘natural’ cutoff, but shall also consider the dependency of the results on the cutoff, by presenting results with cutoff at 0.01, i.e. at a value of order of the slow roll parameters.

Examples

- First, consider the comparison between a trivial Gaussian model (M_1) for which $\lambda_\alpha = 0$ and a DBI-like model (M_2) with free parameter $\lambda_B \neq 0$.

Adopting the Jeffries prior, the Bayes factor between these models is

$$K_{21} = \frac{1}{I} \exp\left(\frac{\chi^2(\lambda_B = 0)}{2}\right) \int_{-10^4}^{10^4} \exp\left(-\frac{\chi^2(\lambda_B)}{2}\right) \frac{d\lambda_B}{|\lambda_B|}. \quad (6.2)$$

The normalization factor I formally satisfies $I = \int_{-10^4}^{10^4} d\lambda_B / |\lambda_B|$, although it is regularized as described in footnote 10. We find $\ln K_{21} \approx 0.63$ which gives no preference for either model. Note that there is an ‘Ockham’s razor’ penalty implicit in (6.2), because the parameter λ_B is allowed to float over a relatively large interval. For a flat prior this Ockham penalty strongly disfavours the DBI model, producing $\ln K_{21} = -4.23$. We conclude that the data are not sufficient to overcome the ambiguity in specifying a prior.

The Bayes factor K_{21} is only one of a number of metrics which can be used to assess goodness of fit. Another is the Akaike ‘information criterion’, defined by $\text{AIC} = \chi_{\text{MLE}}^2 + 2k$, where k measures the number of parameters in the model and is a proxy for the ‘Ockham razor’ penalty of Eq. (6.1). The model with smallest AIC is preferred. We find $\text{AIC}_1 - \text{AIC}_2 = -1.56$, which implies a preference for the trivial Gaussian model M_1 in comparison to a model with nonzero λ_B . The same preference is found if we allow any other single Lagrangian parameter to be nonzero.

- Next, consider a third DBI-like model M_3 in which the two parameters λ_B and λ_D (or, equivalently, the parameters c_s and \tilde{c}_3 in the notation of §5) are allowed to float. We find

		Jeffries prior	
	AIC difference	$ \ln K $, cutoff=1	$ \ln K $, cutoff=0.01
M_1 vs. M_3	-0.94	1.26	0.75
M_2 vs. M_3	0.62	0.64	0.38

The Akaike information criterion prefers M_1 to M_3 , but M_3 to M_2 . Therefore the trivial Gaussian model M_1 is preferred overall, but if we discard this option then the information criterion prefers a two-parameter fit (M_3) to a single-parameter fit (M_2).

The Bayes factors are inconclusive, but it could be argued that they show a weak preference for the opposite conclusion. We compute the Bayes factor using two different choices for the regularization of the Jeffries prior; see footnote 10 on p. 22. A *smaller* cutoff increases the weight of probability for the λ_α to be near zero, and therefore *decreases* the probability that the model generates an observable signature. As we increase the lower limit for the parameters λ_α to the ‘natural’ level $\lambda_\alpha = 1$, the Bayes factor does not strongly discriminate between a two- or three-parameter fit. However, it does marginally begin to disfavour a two-parameter fit (M_3) compared to the trivial model (M_1). Therefore it appears that a fit for a DBI-like model using more than two parameters becomes mildly in tension with the data for ‘natural’ choices of the dimensionless scales λ_α .

The apparent discrepancy with the Akaike information criterion should be ascribed to a stronger ‘Ockham’ or complexity penalty in the Bayes factor. The information criterion down-weights each model by a fixed amount depending on the number of parameters, whereas the Bayes factor attempts to account for the increased volume of parameter space which becomes available. For example, using a flat prior instead of the Jeffries prior very strongly disfavours the models M_2 and M_3 .

7 Discussion and conclusions

The availability of high-quality maps of the CMB temperature anisotropy from the WMAP and Planck missions means that it has become feasible to search for primordial three-point correlations. Such correlations are typically predicted by any scenario in which the fluctuations have an inflationary origin, due to microphysical three-body interactions among the light, active degrees of freedom of the inflationary epoch. If detected, their precise form could provide decisive evidence in favour of the inflationary hypothesis.

Unfortunately, due to issues of computational complexity, it is not yet possible to perform a blind search for these primordial three-point correlations. Instead, we must search for signals which we have some prior reason to believe may be present in the data. Therefore the amount of information we manage to extract depends on which signals we choose to look for.

In this paper we have made a systematic search of the 9-year WMAP data for correlations which could be produced in a very general model of single-field inflation, under the assumption that the background evolution is smooth, yielding corresponding smooth and nearly scale-invariant correlation functions. This excludes models which contain sharp features or oscillations [52–61]. It also excludes models in which significant three-point correlations are generated by differences of evolution between regions of the universe separated by super-Hubble distances. Correlations generated by this mechanism are generally most significant in the ‘squeezed’ or soft limit, where the correlation is between fluctuations on very disparate scales. Such correlations have been disfavoured by analysis of the Planck2013 data release [1]. By comparison, the 9-year WMAP data achieve a smaller signal-to-noise for such configurations. The difference between the 9-year WMAP and Planck2013 datasets is less pronounced for the momentum configurations which we probe, with for example 1σ error bars on $f_{\text{NL}}^{\text{equil}}$ improving from 117 to 75.

The essential steps of our analysis were assembled in §§2–4. We begin with an effective field theory which parametrizes the unknown details of three-body interactions between inflaton fluctuations, but preserves nonlinearly realized Lorentz invariance. The effective theory is agnostic regarding the physical mechanism which underlies inflation. We compute the bispectrum generated by each operator in the effective theory, and break these into principal components using a Fisher-matrix approach. The amplitude of each principal component is recovered from the data, after which the results can be translated into constraints on the mass scales which appeared in the original effective theory. We find that no significant deviation from Gaussianity has been detected in any region of the inflationary parameter space. This conclusion is consistent with previous analyses of the 9-year WMAP and Planck2013 datasets.

Our principal components are similar to those obtained by Byun & Bean, who forecast the constraints which could be obtained from a Planck-like survey [39]. We find that the best-constrained principal direction exhibits similarities to (in order) the flattened, orthogonal and ‘Galileon’ templates. A fourth principal direction is more complex, but at best weakly constrained.

The large space of models which fit into the class of single-field scenarios invites attempts to identify best-fitting regions. To approach this problem we use the framework of Bayesian model comparison. The results are at best weakly significant, but tend to disfavour models with more parameters when compared to simpler cases with zero or one parameter. This is not surprising given that the amplitude of each principal direction is consistent with zero. However, it should be borne in mind that our analysis is restricted to smooth and nearly scale-invariant bispectra. It is possible that a significant signal of a different type is hidden in the data. In some cases, n -point functions of this type can be described within the framework of effective field theory [15]. The analysis developed in §§2–4 could be applied immediately to such scenarios given a suitable choice of basis functions \mathcal{R}_n .

Acknowledgements

It is a pleasure to thank Andrew Liddle for many helpful discussions and useful comments on an advance draft of this manuscript. We are also grateful to Raquel Ribeiro and Sébastien Renaux-Petel for discussions which helped motivate the approach we developed in this work.

Some numerical presented in this paper were obtained using the COSMOS supercomputer, which is funded by STFC, HEFCE and SGL. Other numerical computations were carried out on the Sciama High Performance Compute (HPC) cluster which is supported by the ICG, SEPNet and the University of Portsmouth. DR and DS acknowledge support from the Science and Technology Facilities Council [grant number ST/I000976/1]. GA is supported by STFC grant ST/I506029/1. DS also acknowledges support from the Leverhulme Trust. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007–2013) / ERC Grant Agreement No. [308082].

A Three-point functions for the EFT operators

In this Appendix we briefly recapitulate the calculation of the three-point functions corresponding to each operator in the effective Lagrangian (2.4). The principal tool is Schwinger’s formulation of ‘in–in’ expectation values and the corresponding expansion into diagrams due to Keldysh [62–65]. The technique was applied to general relativity by Jordan, who used it to study the effective equations of motion obtained by integrating out quantum fluctuations [66]. It was imported into cosmology by Calzetta & Hu [67] and applied to inflation by Maldacena and subsequent authors [26, 30, 68].

In–in calculations.—The objective is to calculate the expectation value of a given operator \mathcal{O} at some time t_* , given that the system develops from a specified state (the ‘in’-state) at very early times. This expectation value is

$$\langle \mathcal{O} \rangle_* \equiv \langle \text{in} | \mathcal{O}(t_*) | \text{in} \rangle, \quad (\text{A.1})$$

where the subscript ‘*’ is used to denote evaluation of \mathcal{O} at time t_* . In the present case, \mathcal{O} will correspond to a product of field operators evaluated at the same time but at distinct spatial positions.

Inserting a complete set of intermediate states labelled by the three-dimensional field configuration $\pi(\mathbf{x}, T)$ at some arbitrary time $T > t_*$, we conclude

$$\langle \mathcal{O} \rangle_* = \int [d\pi(\mathbf{x}, T)] \langle \text{in} | \pi(\mathbf{x}, T) \rangle \langle \pi(\mathbf{x}, T) | \mathcal{O}(t_*) | \text{in} \rangle, \quad (\text{A.2})$$

where the measure $[d\pi(\mathbf{x}, T)]$ denotes integration over all field configurations. Each overlap in (A.2) can be written as a conventional Feynman path integral, with the integration running over all field *histories* $\pi(\mathbf{x}, t)$ which are consistent with the in-state $|\text{in}\rangle$ in the far past, and which coincide with the configuration $\pi(\mathbf{x}, T)$ at time T . The result is

$$\langle \mathcal{O}(t_*) \rangle = \int [d\pi_+ d\pi_-] \mathcal{O}(t_*) \exp \left[iS(\pi_+) - iS(\pi_-) \right] \delta[\pi_+(T) - \pi_-(T)], \quad (\text{A.3})$$

with the independent integrations π_+ , π_- running over field histories which are compatible with the in-state but are unrestricted at late times. Eq. (A.3) admits an expansion into diagrams in which the Green's functions connecting only '+' or only '-' fields obey the usual Feynman boundary conditions, but are augmented by Green's functions which mix the '+' and '-' labels and whose boundary conditions are determined by the δ -function. For further details, see Ref. [30].

Mode functions.—It was explained in §3 that we approximate the mode functions as Hankel functions of order $3/2$. Analytically, this corresponds to building Green's functions from the mode function

$$u(\tau, \mathbf{k}) = \frac{iH}{\sqrt{4\epsilon\tilde{c}_s}k^3} (1 + ik\tilde{c}_s\tau) e^{-ik\tilde{c}_s\tau}, \quad (\text{A.4})$$

and its complex conjugate. In this formula, $\tau = -\int_t^\infty dt'/a(t')$ is the conformal time and \tilde{c}_s is a 'generalized' speed of sound. In a model without fourth-derivative kinetic terms this will usually be the phase velocity, determined from the ratio of coefficients of the spatial and temporal kinetic terms. In other cases it may bear less relation to what would normally be thought of as a phase velocity. Our notation coincides with that of Refs. [13, 14], to which we refer for further details; see especially the discussion below Eq. (2.6) in Ref. [14]. In writing Eq. (A.4) we have assumed that the in-state $|\text{in}\rangle$ contains zero particles, corresponding the 'Bunch–Davies' vacuum.

With these choices, the three-point functions corresponding to the EFT operators in (2.4) are:

- $\mathcal{O}_A = -\bar{M}_1^3 (\partial\pi)^2 \partial^2\pi / 2a^4$

$$B_\zeta(k_1, k_2, k_3) \supseteq \frac{1}{16} \bar{M}_1^3 \frac{H_\star^3}{\epsilon^3 c_s^4 \prod_i k_i^3} k_1^2 \mathbf{k}_2 \cdot \mathbf{k}_3 \mathcal{A}_1 + 1 \rightarrow 2 + 1 \rightarrow 3,$$

- $\mathcal{O}_B = -2M_2^4 \dot{\pi} (\partial\pi)^2 / a^2$

$$B_\zeta(k_1, k_2, k_3) \supseteq \frac{1}{8} M_2^4 \frac{H_\star^2}{\epsilon^3 c_s^2 \prod_i k_i^3} k_1^2 \mathbf{k}_2 \cdot \mathbf{k}_3 \left(\frac{1}{k_t} + \frac{k_2 + k_3}{k_t^2} + \frac{2k_2 k_3}{k_t^3} \right) + 1 \rightarrow 2 + 1 \rightarrow 3,$$

- $\mathcal{O}_C = -\bar{M}_2^2 [H(\partial^2\pi)(\partial\pi)^2/2 + \dot{\pi}\partial^2\partial_j\pi\partial_j\pi] / a^4$

$$B_\zeta(k_1, k_2, k_3) \supseteq \frac{1}{16} \bar{M}_2^2 \frac{H_\star^4}{\epsilon^3 c_s^4 \prod_i k_i^3} k_1^2 \mathbf{k}_2 \cdot \mathbf{k}_3 [\mathcal{A}_1 + (k_2^2 + k_3^2) \mathcal{A}_2] + 1 \rightarrow 2 + 1 \rightarrow 3,$$

- $\mathcal{O}_D = -4M_3^4 \dot{\pi}^3/3$

$$B_\zeta(k_1, k_2, k_3) \supseteq \frac{1}{2} M_3^4 \frac{H_\star^2}{\epsilon^3 \prod_i k_i} \frac{1}{k_t^3},$$

- $\mathcal{O}_E = -\bar{M}_3^2 [H(\partial\pi)^2 \partial^2 \pi + \dot{\pi} \partial^2 \partial_j \pi \partial_j \pi] / a^4$

$$B_\zeta(k_1, k_2, k_3) \supseteq \frac{1}{8} \bar{M}_3^2 \frac{H_\star^4}{\epsilon^3 c_s^4 \prod_i k_i^3} k_1^2 \mathbf{k}_2 \cdot \mathbf{k}_3 \left(\mathcal{A}_1 + \frac{k_2^2 + k_3^2}{2} \mathcal{A}_2 \right) + 1 \rightarrow 2 + 1 \rightarrow 3,$$

- $\mathcal{O}_F = -2\bar{M}_4^3 \dot{\pi}^2 \partial^2 \pi / 3a^2$

$$B_\zeta(k_1, k_2, k_3) \supseteq \frac{1}{2} \bar{M}_4^3 \frac{H_\star^3}{\epsilon^3 c_s^2 \prod_i k_i} \frac{1}{k_t^3},$$

- $\mathcal{O}_G = \bar{M}_5^2 \dot{\pi} (\partial^2 \pi)^2 / 3a^4$

$$B_\zeta(k_1, k_2, k_3) \supseteq -\frac{1}{8} \bar{M}_5^2 \frac{H_\star^4}{\epsilon^3 c_s^4 \prod_i k_i} \frac{1}{k_t^3} \left(3 + \frac{4k_\alpha^2}{k_t^2} \right),$$

- $\mathcal{O}_H = \bar{M}_6^2 \dot{\pi} (\partial_i \partial_j \pi)^2 / 3a^4$

$$B_\zeta(k_1, k_2, k_3) \supseteq -\frac{1}{24} \bar{M}_6^2 \frac{H_\star^4}{\epsilon^3 c_s^4 \prod_i k_i^3} k_1^2 (\mathbf{k}_2 \cdot \mathbf{k}_3)^2 \mathcal{A}_2 + 1 \rightarrow 2 + 1 \rightarrow 3,$$

- $\mathcal{O}_I = -\bar{M}_7 (\partial^2 \pi)^3 / 3! a^6$

$$B_\zeta(k_1, k_2, k_3) \supseteq \frac{1}{4} \bar{M}_7 \frac{H_\star^5}{\epsilon^3 c_s^6 \prod_i k_i} \mathcal{A}_3,$$

- $\mathcal{O}_J = -\bar{M}_8 \partial^2 \pi (\partial_j \partial_k \pi)^2 / 3! a^6$

$$B_\zeta(k_1, k_2, k_3) \supseteq \frac{1}{12} \bar{M}_8 \frac{H_\star^5}{\epsilon^3 c_s^6 \prod_i k_i^3} k_1^2 (\mathbf{k}_2 \cdot \mathbf{k}_3)^2 \mathcal{A}_3 + 1 \rightarrow 2 + 1 \rightarrow 3,$$

- $\mathcal{O}_K = -\bar{M}_9 \partial_i \partial_j \pi \partial_j \partial_k \pi \partial_k \partial_i \pi / 3! a^6$

$$B_\zeta(k_1, k_2, k_3) \supseteq \frac{1}{4} \bar{M}_9 \frac{H_\star^5}{\epsilon^3 c_s^6 \prod_i k_i^3} (\mathbf{k}_1 \cdot \mathbf{k}_2)(\mathbf{k}_1 \cdot \mathbf{k}_3)(\mathbf{k}_2 \cdot \mathbf{k}_3) \mathcal{A}_3,$$

where

$$\mathcal{A}_1 = \left(\frac{1}{k_t} + \frac{k_\alpha^2}{k_t^3} + \frac{3k_\beta^3}{k_t^4} \right), \quad \mathcal{A}_2 = \left(\frac{1}{k_t^3} + \frac{3(k_2 + k_3)}{k_t^4} + \frac{12k_2 k_3}{k_t^5} \right), \quad \mathcal{A}_3 = \left(\frac{1}{k_t^3} + \frac{3k_\alpha^2}{k_t^5} + \frac{15k_\beta^3}{k_t^6} \right),$$

with $k_t = k_1 + k_2 + k_3$, $k_\alpha^2 = k_1 k_2 + k_1 k_3 + k_2 k_3$, and $k_\beta^3 = k_1 k_2 k_3$.

References

- [1] **Planck Collaboration** Collaboration, P. Ade *et al.*, *Planck 2013 Results. XXIV. Constraints on primordial non-Gaussianity*, [arXiv:1303.5084](#).
- [2] **WMAP Collaboration** Collaboration, G. Hinshaw *et al.*, *Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results*, [arXiv:1212.5226](#).
- [3] C. Cheung, P. Creminelli, A. L. Fitzpatrick, J. Kaplan, and L. Senatore, *The Effective Field Theory of Inflation*, *JHEP* **0803** (2008) 014, [[arXiv:0709.0293](#)].
- [4] S. Weinberg, *Adiabatic modes in cosmology*, *Phys.Rev.* **D67** (2003) 123504, [[astro-ph/0302326](#)].
- [5] S. Weinberg, *Can non-adiabatic perturbations arise after single-field inflation?*, *Phys.Rev.* **D70** (2004) 043541, [[astro-ph/0401313](#)].
- [6] S. Weinberg, *Must cosmological perturbations remain non-adiabatic after multi-field inflation?*, *Phys.Rev.* **D70** (2004) 083522, [[astro-ph/0405397](#)].
- [7] J. Meyers and N. Sivanandam, *Non-Gaussianities in Multifield Inflation: Superhorizon Evolution, Adiabaticity, and the Fate of f_{nl}*, *Phys.Rev.* **D83** (2011) 103517, [[arXiv:1011.4934](#)].
- [8] J. Elliston, D. J. Mulryne, D. Seery, and R. Tavakol, *Evolution of f_{NL} to the adiabatic limit*, *JCAP* **1111** (2011) 005, [[arXiv:1106.2153](#)].
- [9] P. Creminelli, M. A. Luty, A. Nicolis, and L. Senatore, *Starting the Universe: Stable Violation of the Null Energy Condition and Non-standard Cosmologies*, *JHEP* **0612** (2006) 080, [[hep-th/0606090](#)].
- [10] C. Cheung, A. L. Fitzpatrick, J. Kaplan, and L. Senatore, *On the consistency relation of the 3-point function in single field inflation*, *JCAP* **0802** (2008) 021, [[arXiv:0709.0295](#)].
- [11] S. Weinberg, *Effective Field Theory for Inflation*, *Phys.Rev.* **D77** (2008) 123541, [[arXiv:0804.4291](#)].
- [12] L. Senatore, K. M. Smith, and M. Zaldarriaga, *Non-Gaussianities in Single Field Inflation and their Optimal Limits from the WMAP 5-year Data*, *JCAP* **1001** (2010) 028, [[arXiv:0905.3746](#)].
- [13] N. Bartolo, M. Fasiello, S. Matarrese, and A. Riotto, *Large non-Gaussianities in the Effective Field Theory Approach to Single-Field Inflation: the Bispectrum*, *JCAP* **1008** (2010) 008, [[arXiv:1004.0893](#)].
- [14] N. Bartolo, M. Fasiello, S. Matarrese, and A. Riotto, *Large non-Gaussianities in the Effective Field Theory Approach to Single-Field Inflation: the Trispectrum*, *JCAP* **1009** (2010) 035, [[arXiv:1006.5411](#)].
- [15] N. Bartolo, D. Cannone, and S. Matarrese, *The Effective Field Theory of Inflation Models with Sharp Features*, *JCAP* **1310** (2013) 038, [[arXiv:1307.3483](#)].
- [16] P. Adshead, W. Hu, and V. Miranda, *Bispectrum in Single-Field Inflation Beyond Slow-Roll*, *Phys.Rev.* **D88** (2013) 023507, [[arXiv:1303.7004](#)].
- [17] M. Alishahiha, E. Silverstein, and D. Tong, *DBI in the sky*, *Phys.Rev.* **D70** (2004) 123505, [[hep-th/0404084](#)].
- [18] N. Arkani-Hamed, P. Creminelli, S. Mukohyama, and M. Zaldarriaga, *Ghost inflation*, *JCAP* **0404** (2004) 001, [[hep-th/0312100](#)].
- [19] M. Dias, R. H. Ribeiro, and D. Seery, *The $\hat{I}^{\prime}N$ formula is the dynamical renormalization group*, *JCAP* **1310** (2013) 062, [[arXiv:1210.7800](#)].
- [20] V. Assassi, D. Baumann, and D. Green, *Symmetries and Loops in Inflation*, *JHEP* **1302** (2013) 151, [[arXiv:1210.7792](#)].

- [21] L. Senatore and M. Zaldarriaga, *On Loops in Inflation*, *JHEP* **1012** (2010) 008, [[arXiv:0912.2734](#)].
- [22] L. Senatore and M. Zaldarriaga, *A Naturally Large Four-Point Function in Single Field Inflation*, *JCAP* **1101** (2011) 003, [[arXiv:1004.1201](#)].
- [23] L. Senatore and M. Zaldarriaga, *The Effective Field Theory of Multifield Inflation*, *JHEP* **1204** (2012) 024, [[arXiv:1009.2093](#)].
- [24] D. Baumann and D. Green, *Equilateral Non-Gaussianity and New Physics on the Horizon*, *JCAP* **1109** (2011) 014, [[arXiv:1102.5343](#)].
- [25] D. Baumann, L. Senatore, and M. Zaldarriaga, *Scale-Invariance and the Strong Coupling Problem*, *JCAP* **1105** (2011) 004, [[arXiv:1101.3320](#)].
- [26] J. M. Maldacena, *Non-Gaussian features of primordial fluctuations in single field inflationary models*, *JHEP* **05** (2003) 013, [[astro-ph/0210603](#)].
- [27] P. Creminelli, *On non-Gaussianities in single-field inflation*, *JCAP* **0310** (2003) 003, [[astro-ph/0306122](#)].
- [28] D. Seery and J. E. Lidsey, *Primordial non-gaussianities in single field inflation*, *JCAP* **0506** (2005) 003, [[astro-ph/0503692](#)].
- [29] D. Seery and J. E. Lidsey, *Primordial non-gaussianities from multiple-field inflation*, *JCAP* **0509** (2005) 011, [[astro-ph/0506056](#)].
- [30] S. Weinberg, *Quantum contributions to cosmological correlations*, *Phys.Rev.* **D72** (2005) 043514, [[hep-th/0506236](#)].
- [31] X. Chen, M.-x. Huang, S. Kachru, and G. Shiu, *Observational signatures and non-Gaussianities of general single field inflation*, *JCAP* **0701** (2007) 002, [[hep-th/0605045](#)].
- [32] C. Burrage, R. H. Ribeiro, and D. Seery, *Large slow-roll corrections to the bispectrum of noncanonical inflation*, *JCAP* **1107** (2011) 032, [[arXiv:1103.4126](#)].
- [33] J. Elliston, D. Seery, and R. Tavakol, *The inflationary bispectrum with curved field-space*, *JCAP* **1211** (2012) 060, [[arXiv:1208.6011](#)].
- [34] J. Fergusson, M. Liguori, and E. Shellard, *General CMB and Primordial Bispectrum Estimation I: Mode Expansion, Map-Making and Measures of f_{NL}* , *Phys.Rev.* **D82** (2010) 023502, [[arXiv:0912.5516](#)].
- [35] J. Fergusson, M. Liguori, and E. Shellard, *The CMB Bispectrum*, *JCAP* **1212** (2012) 032, [[arXiv:1006.1642](#)].
- [36] J. Fergusson, D. Regan, and E. Shellard, *Rapid Separable Analysis of Higher Order Correlators in Large Scale Structure*, *Phys.Rev.* **D86** (2012) 063511, [[arXiv:1008.1730](#)].
- [37] J. Fergusson, D. Regan, and E. Shellard, *Optimal Trispectrum Estimators and WMAP Constraints*, [[arXiv:1012.6039](#)].
- [38] D. Regan, E. Shellard, and J. Fergusson, *General CMB and Primordial Trispectrum Estimation*, *Phys.Rev.* **D82** (2010) 023520, [[arXiv:1004.2915](#)].
- [39] J. Byun and R. Bean, *Non-Gaussian Shape Recognition*, *JCAP* **1309** (2013) 026, [[arXiv:1303.3050](#)].
- [40] T. Battfeld and J. Grieb, *Anatomy of bispectra in general single-field inflation – modal expansions*, *JCAP* **1112** (2011) 003, [[arXiv:1110.1369](#)].
- [41] D. Regan, P. Mukherjee, and D. Seery, *General CMB bispectrum analysis using wavelets and separable modes*, [[arXiv:1302.5631](#)].
- [42] D. Regan, M. Gosenca, and D. Seery, *Constraining the WMAP9 bispectrum and trispectrum with needlets*, [[arXiv:1310.8617](#)].

- [43] J. Fergusson and E. Shellard, *The shape of primordial non-Gaussianity and the CMB bispectrum*, *Phys.Rev.* **D80** (2009) 043510, [[arXiv:0812.3413](#)].
- [44] M. Vallisneri, *Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects*, *Phys.Rev.* **D77** (2008) 042001, [[gr-qc/0703086](#)].
- [45] D. Babich, P. Creminelli, and M. Zaldarriaga, *The Shape of non-Gaussianities*, *JCAP* **0408** (2004) 009, [[astro-ph/0405356](#)].
- [46] P. D. Meerburg, J. P. van der Schaar, and P. S. Corasaniti, *Signatures of Initial State Modifications on Bispectrum Statistics*, *JCAP* **0905** (2009) 018, [[arXiv:0901.4044](#)].
- [47] P. Creminelli, G. D’Amico, M. Musso, J. Norena, and E. Trincherini, *Galilean symmetry in the effective theory of inflation: new shapes of non-Gaussianity*, *JCAP* **1102** (2011) 006, [[arXiv:1011.3004](#)].
- [48] R. H. Ribeiro and D. Seery, *Decoding the bispectrum of single-field inflation*, *JCAP* **1110** (2011) 027, [[arXiv:1108.3839](#)].
- [49] **WMAP** Collaboration, C. Bennett *et al.*, *Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results*, *Astrophys.J.Suppl.* **208** (2013) 20, [[arXiv:1212.5225](#)].
- [50] J. Martin, C. Ringeval, R. Trotta, and V. Vennin, *The Best Inflationary Models After Planck*, [[arXiv:1312.3529](#)].
- [51] R. E. Kass and A. E. Raftery, *Bayes factors*, *Journal of the American Statistical Association* **90** (1995), no. 430 773–795, [<http://amstat.tandfonline.com/doi/pdf/10.1080/01621459.1995.10476572>].
- [52] A. A. Starobinsky, *Spectrum of adiabatic perturbations in the universe when there are singularities in the inflation potential*, *JETP Lett.* **55** (1992) 489–494.
- [53] J. A. Adams, G. G. Ross, and S. Sarkar, *Multiple inflation*, *Nucl.Phys.* **B503** (1997) 405–425, [[hep-ph/9704286](#)].
- [54] J. A. Adams, B. Cresswell, and R. Easther, *Inflationary perturbations from a potential with a step*, *Phys.Rev.* **D64** (2001) 123514, [[astro-ph/0102236](#)].
- [55] G. Hailu and S.-H. H. Tye, *Structures in the Gauge/Gravity Duality Cascade*, *JHEP* **0708** (2007) 009, [[hep-th/0611353](#)].
- [56] R. Bean, X. Chen, G. Hailu, S.-H. H. Tye, and J. Xu, *Duality Cascade in Brane Inflation*, *JCAP* **0803** (2008) 026, [[arXiv:0802.0491](#)].
- [57] A. Achucarro, J.-O. Gong, S. Hardeman, G. A. Palma, and S. P. Patil, *Features of heavy physics in the CMB power spectrum*, *JCAP* **1101** (2011) 030, [[arXiv:1010.3693](#)].
- [58] M. Joy, V. Sahni, and A. A. Starobinsky, *A New Universal Local Feature in the Inflationary Perturbation Spectrum*, *Phys.Rev.* **D77** (2008) 023514, [[arXiv:0711.1585](#)].
- [59] S. Hotchkiss and S. Sarkar, *Non-Gaussianity from violation of slow-roll in multiple inflation*, *JCAP* **1005** (2010) 024, [[arXiv:0910.3373](#)].
- [60] M. Nakashima, R. Saito, Y.-i. Takamizu, and J. Yokoyama, *The effect of varying sound velocity on primordial curvature perturbations*, *Prog.Theor.Phys.* **125** (2011) 1035–1052, [[arXiv:1009.4394](#)].
- [61] P. Adshead, W. Hu, C. Dvorkin, and H. V. Peiris, *Fast Computation of Bispectrum Features with Generalized Slow Roll*, *Phys.Rev.* **D84** (2011) 043519, [[arXiv:1102.3435](#)].
- [62] J. S. Schwinger, *Brownian motion of a quantum oscillator*, *J.Math.Phys.* **2** (1961) 407–432.

- [63] P. M. Bakshi and K. T. Mahanthappa, *Expectation value formalism in quantum field theory. 1.*, *J.Math.Phys.* **4** (1963) 1–11.
- [64] P. M. Bakshi and K. T. Mahanthappa, *Expectation value formalism in quantum field theory. 2.*, *J.Math.Phys.* **4** (1963) 12–16.
- [65] L. Keldysh, *Diagram technique for nonequilibrium processes*, *Zh.Eksp.Teor.Fiz.* **47** (1964) 1515–1527.
- [66] R. Jordan, *Effective Field Equations for Expectation Values*, *Phys.Rev.* **D33** (1986) 444–454.
- [67] E. Calzetta and B. Hu, *Closed Time Path Functional Formalism in Curved Space-Time: Application to Cosmological Back Reaction Problems*, *Phys.Rev.* **D35** (1987) 495.
- [68] S. Weinberg, *Quantum contributions to cosmological correlations. II. Can these corrections become large?*, *Phys.Rev.* **D74** (2006) 023508, [[hep-th/0605244](#)].

	λ_A	λ_B	λ_C	λ_D	λ_E	λ_F	λ_G	λ_H	λ_I	λ_J	λ_K
λ_1	-0.173206	-0.193711	-0.221619	-0.260354	-0.203464	-0.260354	-0.256692	-0.507151	-0.24732	-0.446539	0.350338
λ_2	-0.240889	-0.222667	-0.200172	-0.163445	-0.215441	-0.163445	-0.165617	0.060200	-0.171386	0.048317	-0.830493
λ_3	0.269654	0.191121	0.050094	-0.064110	0.132429	-0.064110	0.006499	-0.782966	0.226131	0.379275	-0.233301
λ_4	-0.219333	-0.233381	-0.193386	-0.279000	-0.203131	-0.279000	-0.120479	0.176132	0.481877	0.541142	0.304148

Table 3: Linear combinations of the λ_α parameters which can be constrained using the WMAP9 bispectrum data. The unrotated parameters are labelled $\lambda_A, \lambda_B, \dots$, and correspond to those defined in Table 1 in terms of the EFT mass scales. The rotated parameters are labelled $\lambda_1, \lambda_2, \dots$